Optimum Perfect Universal Steganography of Finite Memoryless Sources

Félix Balado and David Haughton

arXiv:1410.2659v2 [cs.IT] 14 Oct 2014

***Abstract*—A theoretical analysis and a practical solution to the fundamental problem of optimum perfect universal steganography of finite memoryless sources with a passive warden are provided. The theoretical analysis rests on the fact that Slepian's Variant I permutation coding implements first-order perfect universal steganography with maximum embedding rate. The practical solution underlines the duality between perfect steganography with optimum embedding rate and lossless source coding with optimum compression rate, since it is shown that permutation coding can be implemented by means of adaptive arithmetic coding. A distance constraint between host signal and information-carrying signal must be observed in universal steganography, and thus an optimum tradeoff between embedding rate and embedding distortion must be achieved. Partitioned permutation coding is shown to be the solution to this problem. A practical implementation of partitioned permutation coding is given that performs close to an unattainable upper bound on the rate-distortion function of the problem.**

***Index Terms*—Steganography, source coding, permutation coding, arithmetic coding, rate-distortion.**

I. Introduction

DIGITAL data hiding refers to coding techniques which aim at embedding information within digital discrete-time host signals [1]. In short, steganography is a special data hiding scenario in which undetectability of the embedded information is paramount —unobtrusiveness, rather than undetectability, suffices in general data hiding. In the steganographic problem, a "man in the middle" (warden) performs a detection test on signals sent between two parties in order to determine whether they carry hidden information or not. In the scenario considered here, the warden does not alter the tested signals, which thus arrive unmodified at the decoder (passive warden). Besides circumventing detection by the warden, the encoder also wishes to maximize the steganographic embedding rate, which is the amount of bits per host element that can be conveyed to the decoder through the covert channel created by modifying the host in order to embed information.

An important landmark in steganography research was the realization of the existence of an inextricable connection between steganography and source coding. Apart from this connection being already explored in the influential early work by Cachin [2], Anderson and Petitcolas gave, around the same time, an intuitive rationale for this link [3, Section VI-A]. These authors pointed out that if we would have lossless source coding with optimum compression rate for real-world signals (i.e., ideal compression for signals such as digital images), then these signals would have to be dense in the space of the optimum source code. Thus, decompressing any arbitrary message from this space would always render an innocuous-looking real-world signal. In other words, a lossless compression algorithm with optimum compression rate could also work as a perfect (undetectable) steganographic algorithm with optimum embedding rate, just by using decompression to encode information-carrying signals and compression to decode that information.

These observations about the duality between perfect steganography with a passive warden and lossless source coding —in fact, a particular manifestation of the duality between channel coding and source coding first observed by Shannon [4]— implied that practical steganographic algorithms had to be intimately related to practical compression algorithms. A number of early works partially succeeded in translating this fundamental relationship into well-founded steganographic methods. The first such work was Cachin's [2]. Later on, Sallee added another important piece to the puzzle with model-based steganography [5]. Unfortunately, steganography research has gradually drifted away from these seminal contributions, and steganography has become a subject largely disconnected from the principles of source coding. It is also disappointing to see that recent practical steganographic methods have largely ignored relevant information-theoretical results such as [6] or [7]. Hence, one cannot help but conclude that the current state of affairs in the steganography field is unsatisfactory. The recent trend in steganography can be summarized as "stress-testing of heuristic algorithms by means of machine-learning detectors." However, it would be hard to contend that this line of attack —detached from the theoretical foundations of steganography— can provide answers to the deeper questions of the discipline, notwithstanding the fact that partial successes may be in fact achieved through empirical approaches.

Here we present a contribution which we believe fills an important gap in the field: the solution to a canonical problem in steganography, showing its connection with optimum source coding and with existing theoretical results on steganographic systems. In order to understand the necessity of studying the problem that we will address in this paper, we must first cast a brief look at some key elements in the history of compression algorithms —a relevant exercise due to the fact that source coding has traveled a longer and arguably more fruitful journey as a discipline than steganography. A canonical

F. Balado and D. Haughton are with the School of Computer Science and Informatics, University College Dublin, Belfield Campus, Dublin 4, Ireland (e-mails: felix@ucd.ie, david.haughton@ucdconnect.ie). This research was conducted with the financial support of Science Foundation Ireland under Grant Number 09/RFP/CMS2212. Preliminary parts of this work were presented at the IEEE ICASSP Conference (May 2013), IEEE WIFS Conference (November 2013), and the WIO Conference (July 2014).

problem in lossless source coding is how to optimally compress a signal whose samples are independently drawn from a known distribution. The theoretical limit to this problem and a prototype of its practical solution were known since Shannon (Shannon-Fano-Elias coding). Yet, a complete practical solution only emerged in 1976, when Pasco [8] and Rissanen [9] independently and almost simultaneously developed arithmetic coding: a lossless source coding algorithm whose compression rate asymptotically approaches that of the entropy of the signal to be compressed. Arithmetic coding is a cornerstone of modern compression methods despite two practical issues. To start with, real-world signals are not memoryless, but rather exhibit statistical dependencies. This problem is dealt with in practical compression by means of decorrelation techniques that approximate the statistical independence hypothesis. A further difficulty is that the distribution of the signal to be compressed is commonly unknown. The pragmatic answer to this issue is that the empirical distribution of the signal can always act as a proxy for its underlying unknown distribution (universal compression). Consequently, the fundamental problem of lossless universal source coding of memoryless signals with near-optimum compression rate can be solved, in practice, by means of arithmetic coding plus adaptive (on the fly) estimation of the first-order statistics of the digital signal to be compressed [10].

Let us now turn our attention to steganography. The celebrated Cachin's criterion [2] tells us that a perfectly secure steganographic system is implemented by an encoder that preserves the distribution of the host signal, because in this way optimum detection by the warden will be foiled. The implementation of this central criterion raises a crucial issue which —not accidentally— we have already found in our short review of source coding: what should the encoder do if the distribution of the host signal is not known? As noted by Cachin himself by drawing the parallel with source coding, universal steganography must necessarily be the matter-of-fact approach to implementing perfect steganography: it is the empirical distribution of the host that should be preserved [2]. Thus, a canonical problem in steganography —which echoes its source coding counterpart— is how to undertake perfect universal steganography of memoryless host signals with optimum embedding rate. Clearly, if the host signal is memoryless then perfect universal steganography is achieved by preserving its first-order statistics. However the problem has remained open to this date, despite of the fact that some authors have given important steps toward its solution. Cachin was the first to sketch an answer, relying on a description of a generic universal compressor based on the method of types [2]. However the framework proposed by Cachin suffers from two shortcomings: 1) it does not provide perfect steganography for finite hosts, as it only achieves perfection asymptotically when the size of the host goes to infinity; 2) it assumes that all signals with the same empirical distribution as the host are semantically equivalent to it.

The first shortcoming was addressed by Ryabko and Ryabko [11], who proposed a perfect universal steganographic algorithm with optimum embedding rate for finite hosts. However this method also fails to address the second shortcoming, which is the most acute in practice as we will argue next. Recalling the arguments by Anderson and Petitcolas for the duality between steganography and source coding, one must observe that, in their idealized steganographic problem, the encoder produces a semantically meaningful information-carrying signal *ab initio*, relying on an ideal model of the signals that the encoder can produce and without the need of a host signal. Yet, in a universal approach, such as in [2] or in [11], the role of the hypothetical ideal model is played by the actual empirical model of a host signal. A key observation must be made at this point: nothing guarantees that a signal which preserves the empirical model of the host will also be semantically equivalent to it, or even semantically meaningful at all. Hence a signal which is simply bootstrapped from the empirical distribution of the host will generally be suspicious, and thus detectable by the warden. But where does the qualifier "semantically meaningful" belong in signals that are memoryless? We have mentioned that decorrelation is a way to approximate the memoryless assumption for real-world signals. Clearly, not all signals that preserve the first-order statistics of a host signal in the decorrelated domain will map back to semantically meaningful signals in the original domain. Therefore, a distance constraint between host and information-carrying signal (embedding distortion constraint) is essential in universal steganographic approaches. This implies that, unlike the idealized problem, practical universal steganography must be a problem of coding with side information at the encoder —like virtually all other data hiding problems.

In regard to the embedding distortion constraint, a very interesting practical approach to the problem of near-perfect steganography of finite memoryless sources with near-optimum embedding rate was given by Sallee [5]. Sallee used the quantized block-DCT domain as a rough approximation to a domain where the host is memoryless, and in which probabilistic models are available. Exploiting these two properties and the steganography-source coding duality, he proposed a scheme called model-based steganography which uses arithmetic coding to achieve near-maximum embedding rate while preserving a first-order model of the host, and which simultaneously limits the embedding distortion incurred by the encoder. Unfortunately, due to its reliance on a theoretical model of the host, model-based steganography is neither universal nor perfect.

In order to obtain a solution to the canonical problem that occupies us in this paper, we will take an approach whose core element is —perhaps paradoxically when considering the duality arguments presented throughout this introduction— not a lossless compression method, but rather Slepian's Variant I permutation coding [12]. As it will be discussed, it is straightforward to see that permutation coding implements perfect universal steganography of finite memoryless hosts with maximum embedding rate. The centrality of permutations to the problem at hand had already been discovered —either explicitly or not— by a number of researchers over the years, most prominently by Ryabko and Ryabko [11]. However, the question has remained undeveloped in terms of answers to essential questions: what is the optimum embedding rate for a given embedding distortion constraint, and how does one

implement a near-optimum algorithm in practice? Finally, as we have already mentioned, an approach to steganography informed by permutation coding may appear to be unrelated to compression, despite the fact that the duality between these two fields runs deep. In this respect, we believe that a relevant contribution of this work is an implementation of permutation coding based on near-optimum lossless compression.

This paper is organized as follows. Section II introduces the notational conventions and the basic framework used throughout the paper, and Section III reviews prior work on perfect steganography of finite memoryless hosts. Section IV is devoted to a theoretical analysis of the embedding distortion and the embedding rate of permutation coding in a steganographic context. Next, Section V addresses the two main practical issues found in the implementation of permutation coding in steganography. Theoretical and empirical results are compared in Section VI, and, lastly, Section VII draws the conclusions of this work.

II. Notation and Framework

Boldface lowercase Roman letters are column vectors. The i-th element of vector $\mathbf{a}$ is a_i, although for notational clarity we will also use $(\mathbf{a})_i$ whenever convenient. The special symbols $\mathbf{1}$ and $\mathbf{0}$ are the all-ones vector and the null vector, respectively. Capital Greek letters denote matrices; the entry at row i and column j of matrix Π is $(\Pi)_{i,j}$. For the sake of keeping standard conventions, the only two exceptions to this convention are the identity matrix I and the exchange matrix J (which will be defined later). $\operatorname{tr}\Pi$ is the trace of the square matrix Π. $(\cdot)^t$ is the vector or matrix transpose operator. $\operatorname{diag}(\mathbf{a})$ is a diagonal matrix with vector $\mathbf{a}$ in its diagonal. The indicator function is defined as $\mathbb{1}_{\{\theta\}} = 1$ if logical expression θ is true, and zero otherwise. The 2-norm of a vector $\mathbf{r}$ is $\|\mathbf{r}\| = \sqrt{\mathbf{r}^t\mathbf{r}}$. The Hamming distance between two n-vectors $\mathbf{r}$ and $\mathbf{s}$ is $\delta(\mathbf{r},\mathbf{s}) = \sum_{i=1}^{n} \mathbb{1}_{\{r_i \neq s_i\}}$. Calligraphic letters are sets; $|\mathcal{V}|$ is the cardinality of set $\mathcal{V}$.

A host sequence is denoted by the discrete-valued n-vector $\mathbf{x} = [x_1, x_2, \ldots, x_n]^t \in \mathcal{V}^n$ where $\mathcal{V} = \{v_1, v_2, \ldots, v_q\} \subset \mathbb{Z}$. We assume that $\mathbf{x} \neq \mathbf{0}$, and also that $\mathbf{v} = [v_1, v_2, \ldots, v_q]^t$ gives the elements of $\mathcal{V}$ in increasing order, that is, $v_1 < v_2 < \cdots < v_q$. The histogram of $\mathbf{x}$ is a vector $\mathbf{h} = [h_1, h_2, \ldots, h_q]^t$ such that $h_k = \sum_{i=1}^{n} \mathbb{1}_{\{v_k = x_i\}}$ for $k = 1, 2, \ldots, q$, and then $n = \mathbf{h}^t\mathbf{1}$; $\mathbf{v}$ is therefore the vector containing the ordered histogram bins. An information-carrying sequence is denoted by $\mathbf{y} = [y_1, y_2, \ldots, y_n]^t$.

Let $\mathcal{S}_n$ be the symmetric group, that is, the group of all permutations of $\{1, 2, \ldots, n\}$. We denote a permutation $\boldsymbol{\sigma} \in \mathcal{S}_n$ by means of a vector $\boldsymbol{\sigma} = [\sigma_1, \sigma_2, \ldots, \sigma_n]^t$ where $\sigma_i \in \{1, 2, \ldots, n\}$ and $\sigma_i \neq \sigma_j$ for all $i \neq j$. This vector can be used in turn to define a permutation matrix $\Pi_{\boldsymbol{\sigma}}$ with entries $(\Pi_{\boldsymbol{\sigma}})_{i,j} = \mathbb{1}_{\{\sigma_i = j\}}$. The rearrangement of $\mathbf{x}$ using $\boldsymbol{\sigma}$ is the vector $\mathbf{y} = \Pi_{\boldsymbol{\sigma}}\,\mathbf{x}$, for which $y_i = x_{\sigma_i}$ for $i = 1, 2, \ldots, n$. Observe that two or more different permutations may lead to the same reordering of the elements of $\mathbf{x}$. For this reason we will follow the convention that a *rearrangement* of $\mathbf{x}$ is a unique ordering of its elements. A special case is the rearrangement of $\mathbf{x}$ in nondecreasing order. This is obtained by means of a permutation $\overrightarrow{\boldsymbol{\sigma}}$ yielding $\overrightarrow{\mathbf{x}} = \Pi_{\overrightarrow{\boldsymbol{\sigma}}}\,\mathbf{x}$ such that $\overrightarrow{x}_1 \leq \overrightarrow{x}_2 \leq \cdots \leq \overrightarrow{x}_n$. The rearrangement of $\mathbf{x}$ in nonincreasing order can be obtained from $\overrightarrow{\mathbf{x}}$ as $\overleftarrow{\mathbf{x}} = \mathrm{J}\overrightarrow{\mathbf{x}}$, where J is the exchange matrix —the permutation matrix with entries $(\mathrm{J})_{i,j} = \mathbb{1}_{\{j = n-i+1\}}$.

Italicized Roman or Greek capital letters represent random variables. The probability mass function (pmf) of a random variable X with support $\mathcal{V}$ is denoted by $p(X = v)$, with $v \in \mathcal{V}$, or simply by $p(v)$ if clear from the context. We will also refer to $\mathbf{p} = [p(X = v_1), p(X = v_2), \ldots, p(X = v_q)]^t$ as the pmf of X, and to $\mathbf{v}$ as its support. The probability of an event ϑ is denoted by $\Pr\{\vartheta\}$. The expectation, variance, and entropy of X are denoted by $\mathrm{E}\{X\}$, $\mathrm{Var}\{X\}$, and $H(X)$, respectively. $I(X;Y)$ is the mutual information between X and Y. Logarithms are base 2 throughout the paper, unless explicitly noted otherwise.

III. Prior Work

As discussed in the introduction, we will be dealing with universal perfect steganography of finite memoryless hosts. If the host is discrete-valued, preserving its first-order statistics is equivalent to preserving its histogram. A number of previous authors more or less explicitly claim that their steganographic algorithms implement histogram preservation. However, on close examination, many allegedly histogram-preserving methods only implement approximations to histogram preservation. We will briefly review next the few steganographic methods that do implement exact histogram preservation, which therefore implement perfect steganography of finite memoryless sources. Among them we find Provos' OutGuess [13] —the first histogram-preserving steganographic algorithm—, Franz's proposal [14], Ryabko and Ryabko's method [11], Kumar and Newman's J3 [15], and Luo and Subbalakshmi's method [16].

Most of these methods (see [13], [14], [15]) are variations of least-significant bit (LSB) steganography. In this heuristic early steganographic method a message is embedded into a host $\mathbf{x}$ by producing an information-carrying signal $\mathbf{y}$ whose elements are $y_i = 2\lfloor x_i/2 \rfloor + b_i$ for $i = 1, 2, \ldots, n$, where $b_i \in \{0, 1\}$ is the message bit embedded in the i-th element of the host. The histogram-preserving methods in [13], [14] and [15] can be seen as using this embedding scheme plus some kind of histogram compensation to remediate the lack of histogram preservation in the baseline technique. These beginnings limit their theoretical and practical outlook from inception. For instance, all of the methods just cited possess a low embedding rate. The report in [13] suggests that the average embedding rate of Outguess lies around 0.31 bits/host element, whereas Franz's method is below 0.20 bits/host element for most of the hosts tested in [14], and J3 offers rates between 0.35 bits/host element and 0.65 bits/*nonzero* host element [15], all clearly below the ceiling embedding rate of 1 bit/host element implemented by LSB steganography.

The notable exception among all previous works in terms of generality and embedding rate is the proposal by Ryabko and Ryabko [11], which realizes the following fundamental observation: any information-carrying vector $\mathbf{y}$ that preserves the histogram of $\mathbf{x}$ must be a rearrangement of $\mathbf{x}$ (rearrangements are called orderings in [11]). This is because histogram

preservation implies that $\sum_{i=1}^{n} \mathbb{1}_{\{v_k=y_i\}} = \sum_{i=1}^{n} \mathbb{1}_{\{v_k=x_i\}}$ for all $k = 1, 2, \ldots, q$, which can only be true if $\mathbf{y} = \Pi_{\boldsymbol{\sigma}}\, \mathbf{x}$ for some permutation $\boldsymbol{\sigma} \in \mathcal{S}_n$. This observation was first made by Franz [14], but Ryabko and Ryabko [11] pursued it further by proposing a method with near optimum embedding rate. However they did not exploit the duality between perfect steganography and lossless compression, which as we will see, leads to a simple solution to the problem considered by them. As for other steganographic algorithms explicitly based on rearrangements, they are clearly limited in terms of their embedding rate, since they do not exploit the whole spectrum of histogram-preserving possibilities. For instance [16], which is simply based on permuting pairs of host elements, can only achieve a ceiling rate of 0.5 bits/host element.

Finally, we have argued in the introduction that enforcing an embedding distortion constraint is essential in universal steganography. Ryabko and Ryabko [11] did not consider any such constraints, and so their method is only near optimum in unconstrained settings. Other histogram-preserving methods [13], [14], [15], [16] do implement embedding distortion control but in ad-hoc manners, which is why they are not amenable to near-optimum rate-distortion tradeoff.

IV. Permutation Codes as Steganographic Codes

After our brief review of prior histogram-preserving methods, let us explore next the implications of taking to its full extent the observation in the previous section about histogram preservation necessarily involving rearrangements of the host. First of all, this observation means that the perfect steganographic codes for finite memoryless hosts are the Variant I permutation codes first described by Slepian [12]. We must mention that Mittelholzer [17] was the first to consider Slepian's permutation modulation as a data hiding tool. However he did not investigate the histogram-preserving case $\mathbf{y} = \Pi_{\boldsymbol{\sigma}}\mathbf{x}$, that mirrors the actual use of permutation codes in channel/source coding [12], [18] —i.e., the case in which all codewords $\mathbf{y}$ are permutations of a base vector $\mathbf{x}$— but rather he studied the non-histogram-preserving case $\mathbf{y} = \mathbf{x} + \Pi_{\boldsymbol{\sigma}}\mathbf{k}$, with $\mathbf{k}$ a secret vector, which is not relevant to our problem.

Let us state next some basic facts and definitions about permutation coding which we will use throughout the paper. The number r of rearrangements of $\mathbf{x}$ only depends on its histogram, and it is given by the following multinomial coefficient:

$$r \triangleq \binom{n}{\mathbf{h}} = \frac{n!}{h_1! h_2! \cdots h_q!}. \tag{1}$$

Hereafter we will consider $\mathcal{S}_{\mathbf{x}} \subset \mathcal{S}_n$ to be any set of permutations leading to the $r = |\mathcal{S}_{\mathbf{x}}|$ rearrangements of $\mathbf{x}$. Let us denote the r rearrangements of the host $\mathbf{x}$ as $\mathbf{y}^{(1)}, \mathbf{y}^{(2)}, \ldots, \mathbf{y}^{(r)}$, which are the codewords that the encoder can produce. The encoder is therefore a function $e(\cdot,\cdot) : \mathcal{V}^n \times \{1, 2, \ldots, r\} \to \mathcal{V}^n$ such that $e(\mathbf{x}, m) = \mathbf{y}^{(m)}$, where m is the message to be embedded. An alternative view, which is convenient to study the embedding distortion, is to see the encoding procedure as embedding an additive *watermark* $\mathbf{w}^{(m)} \triangleq \mathbf{y}^{(m)} - \mathbf{x}$ into the host $\mathbf{x}$. On the other hand the decoder is a function $d(\cdot) : \mathcal{V}^n \to \{1, 2, \ldots, r\}$ that retrieves the message embedded in a received vector. Therefore the decoded message can be put as $\hat{m} = d(\mathbf{y})$, and with a passive warden $d\left(\mathbf{y}^{(m)}\right) = m$. In the remainder we will drop the superindex (m) from $\mathbf{y}^{(m)}$ and $\mathbf{w}^{(m)}$ for clarity of exposition, whenever this is unambiguous from the context. As we will see in Section V-A3, the encoding and decoding functions can also include a symmetric secret key for privacy purposes, but for simplicity we will omit the key from most of our exposition without loss of generality.

The central difference between the use of permutation codes in channel/source coding [12], [18] and in perfect steganography is that $\mathbf{x}$ is a design choice in the former case, but it is a fixed input parameter of the encoder in the latter. For example, in the communications scenario with a white additive Gaussian channel, $\mathbf{x}$ is chosen so as to maximize the minimum distance between pairs of codewords, because in this case maximum likelihood decoding amounts to minimum distance decoding [12]. However in the steganographic scenario $\mathbf{x}$ is the host signal, and, as such, given to the encoder as a fixed input parameter. As we will discuss in Section V, this simple fact is at the root of the two most relevant issues found in the use of permutation coding in steganography: encoder complexity (due to the fact that $\mathbf{x}$ may be a very long vector) and embedding distortion control. In order to address these issues, we firstly need to undertake a detailed theoretical analysis of permutation coding for steganography. Among other things, this analysis will reveal that permutation coding possesses a number of properties that had been theoretically inferred about optimum perfect steganographic methods by previous authors.

A. Embedding Distortion

In this section we will analyze the theoretical embedding distortion of permutation coding, without occupying ourselves for the time being with the practical question of how to control it. However, as we will see in Section V-B, this theoretical analysis will be key in the practical implementation of permutation coding.

1) Squared Euclidean distance: A very useful way to measure the embedding distortion is by means of the squared Euclidean distance between a codeword $\mathbf{y}$ and the host $\mathbf{x}$, namely, the squared 2-norm of the corresponding watermark $\|\mathbf{w}\|^2 = \|\mathbf{y} - \mathbf{x}\|^2$ —which we will also refer to as the power of $\mathbf{w}$. The main reasons for considering this embedding distortion measure are the following ones: 1) it enables direct comparisons with prior research results when normalized by the power of the host, which yields communications-like signal to noise ratios widely used in data hiding (see Section IV-A1a); 2) it is amenable to analysis and, as it will be seen, it provides deep insights about the use of permutation coding in steganography, both in terms of geometry and of rate-distortion properties; and 3) if $\mathbf{x}$ is the product of a decorrelating unitary linear transform, then the watermark power $\|\mathbf{w}\|^2$ is preserved in the original (correlated) domain.

Using the fact that all histogram-preserving codewords $\mathbf{y}$ have the same 2-norm $\|\mathbf{y}\| = \|\mathbf{x}\|$, the power of a histogram-preserving watermark can be put as

$$\|\mathbf{w}\|^2 = 2\left(\|\mathbf{x}\|^2 - \mathbf{x}^t\mathbf{y}\right) = 2\left(\|\mathbf{x}\|^2 - \mathbf{x}^t\Pi_{\boldsymbol{\sigma}}\mathbf{x}\right) \tag{2}$$

for some $\boldsymbol{\sigma} \in \mathcal{S}_{\mathbf{x}}$. Since this distortion is dependent on the message associated to $\boldsymbol{\sigma}$, we will derive next two relevant message-independent embedding distortion measures, which will be seen to completely suffice in order to approximate and/or bound (2) for any $\boldsymbol{\sigma}$.

- *Average watermark power.* If the encoder chooses messages uniformly at random, then the average watermark power is $\overline{\|\mathbf{w}\|^2} \triangleq (1/r)\sum_{m=1}^{r}\|\mathbf{w}^{(m)}\|^2 = 2\left(\|\mathbf{x}\|^2 - (1/r)\sum_{\boldsymbol{\sigma}\in\mathcal{S}_{\mathbf{x}}}\mathbf{x}^t\,\Pi_{\boldsymbol{\sigma}}\mathbf{x}\right)$. Using next expressions (68) and (69) from the Appendix, and observing that $\mathbf{x}^t\mathbf{1}\mathbf{1}^t\mathbf{x} = (\mathbf{x}^t\mathbf{1})^2$, one arrives at

$$\overline{\|\mathbf{w}\|^2} = 2\left(\|\mathbf{x}\|^2 - \frac{1}{n}(\mathbf{x}^t\mathbf{1})^2\right). \tag{3}$$

  Since $(\mathbf{x}^t\mathbf{1})^2 \geq 0$ it holds that

$$\overline{\|\mathbf{w}\|^2} \leq 2\|\mathbf{x}\|^2, \tag{4}$$

  with equality for any zero-sum $\mathbf{x}$. As we will see throughout the paper, the average watermark power plays a pivotal role in the application of permutation coding to steganography. For a start, we verify next that (3) indicates already that permutation coding satisfies fundamental theoretical requirements of perfect steganography. The average watermark power per host element can be put as

$$\frac{1}{n}\overline{\|\mathbf{w}\|^2} = 2\,\sigma_{\mathbf{x}}^2, \tag{5}$$

  where $\sigma_{\mathbf{x}}^2$ is the (biased) sample variance of $\mathbf{x}$. Since the maximum embedding rate is achieved when the encoder is free to generate all rearrangements of $\mathbf{x}$, then (5) is the exact coding analogous of the theoretical result by Comesaña and Pérez-González in [6, page 17] showing that the average quadratic embedding distortion in an unconstrained capacity-achieving perfect steganographic method is

$$\frac{1}{n}\,\mathrm{E}\{\|\boldsymbol{W}\|^2\} = 2\,\mathrm{Var}\{X\}, \tag{6}$$

  where X is a random variable describing an independent and identically distributed (i.i.d.) host, and $\boldsymbol{W}$ is a random n-vector describing a perfect watermark.
- *Maximum watermark power.* It is also desirable to obtain the maximum power of a perfect steganographic watermark, $\|\mathbf{w}\|_{\max}^2 \triangleq \max_{m\in\{1,2,\ldots,r\}}\|\mathbf{w}^{(m)}\|^2$, which is the worst-case embedding distortion associated to a single message. To this end we may use the following rearrangement inequality [19, Chapter 10]:

$$\overrightarrow{\mathbf{r}}^t\overleftarrow{\mathbf{s}} \leq \mathbf{r}^t\mathbf{s}, \tag{7}$$

  which holds for any $\mathbf{r}, \mathbf{s} \in \mathbb{R}^n$. Setting $\mathbf{r} = \mathbf{x}$ and $\mathbf{s} = \mathbf{y}$ and observing that $\mathbf{y}$ is a rearrangement of $\mathbf{x}$ we have from (2) and (7) that

$$\|\mathbf{w}\|_{\max}^2 = 2\left(\|\mathbf{x}\|^2 - \overrightarrow{\mathbf{x}}^t\overleftarrow{\mathbf{x}}\right). \tag{8}$$

  An upper bound on $\|\mathbf{w}\|_{\max}^2$ is not immediately apparent from inspecting (8), because $\overrightarrow{\mathbf{x}}^t\overleftarrow{\mathbf{x}}$ may be negative. Since $\mathrm{J} = \mathrm{J}^t$, in order to bound $\|\mathbf{w}\|_{\max}^2$ from above we can use the Rayleigh-Ritz theorem to write $\overrightarrow{\mathbf{x}}^t\overleftarrow{\mathbf{x}} = \overrightarrow{\mathbf{x}}^t\mathrm{J}\overrightarrow{\mathbf{x}} \geq \lambda_{\min}(\mathrm{J})\|\mathbf{x}\|^2$, where $\lambda_{\min}(\mathrm{J})$ is the minimum eigenvalue of J. By definition, an eigenvalue λ of J and its associated eigenvector $\mathbf{u}$ fulfill $\mathrm{J}\mathbf{u} = \lambda\mathbf{u}$. Multiplying this expression by $\mathbf{u}^t$ one obtains $\mathbf{u}^t\mathrm{J}\mathbf{u} = \lambda\|\mathbf{u}\|^2$; alternatively, multiplying it by $\mathbf{u}^t\mathrm{J}$ one obtains $\|\mathbf{u}\|^2 = \lambda\mathbf{u}^t\mathrm{J}\mathbf{u}$, because the exchange matrix is involutory ($\mathrm{JJ} = \mathrm{I}$). Combining these two equations one sees that $\lambda^2 = 1$, and since $\mathrm{tr}\,\mathrm{J} \in \{0, 1\}$ then $\lambda_{\min}(\mathrm{J}) = -1$ when $n > 1$. Therefore

$$\|\mathbf{w}\|_{\max}^2 \leq 4\|\mathbf{x}\|^2. \tag{9}$$

  As we will see in Section IV-A1c, this inequality can be more directly obtained through geometric arguments; we will discuss when equality occurs in (9) in that section.

Finally, a trivial inequality involving (3) and (8) is

$$\overline{\|\mathbf{w}\|^2} \leq \|\mathbf{w}\|_{\max}^2, \tag{10}$$

in which equality occurs when $\mathbf{x} = v\mathbf{1}$. In this case $\overline{\|\mathbf{w}\|^2} = \|\mathbf{w}\|_{\max}^2 = 0$ and $r = 1$.

a) Power ratios: The embedding distortion expressions (3) and (8) must be normalized using the power of the host in order to be meaningful across different hosts. Thus, the following figures of merit for the theoretical embedding distortion can be put forward:

- Host to average watermark power ratio

$$\underline{\xi} \triangleq \frac{\|\mathbf{x}\|^2}{\overline{\|\mathbf{w}\|^2}}. \tag{11}$$

  A common alternative to $\underline{\xi}$ is the peak host to average watermark power ratio $\underline{\xi}^* \triangleq n(2^b-1)^2/\overline{\|\mathbf{w}\|^2}$, where it is assumed that $\mathbf{x}$ is in the nonnegative orthant and represented using b bits/element.
- Host to maximum watermark power ratio

$$\xi_{\min} \triangleq \frac{\|\mathbf{x}\|^2}{\|\mathbf{w}\|_{\max}^2}.$$

These figures of merit are trivially related as follows:

$$\underline{\xi}^* \geq \underline{\xi} \geq \xi_{\min}. \tag{12}$$

In keeping with standard conventions and where convenient throughout the paper, we will also refer to the ξ ratios in terms of decibels (dB), by which the amount $10\log_{10}\xi$ is understood. For all the convenience of these figures of merit, a word of caution is needed: it is clear that, as a rule, low ξ ratios imply semantic dissimilarity between $\mathbf{x}$ and $\mathbf{y}$ (in practice, of their counterparts in the original correlated domain), but the converse is not always true. A well-known example of this shortcoming —which is due to the simplification of using the squared Euclidean distance to gauge semantic similarity— is the widely used peak signal to noise ratio (PSNR) of which $\underline{\xi}^*$ is a version. According to our discussion, high ξ ratios are a necessary condition to achieve semantic fidelity. However (4) implies that $\underline{\xi} \geq 1/2$ (≈ -3 dB), whereas (9) implies that $\xi_{\min} \geq 1/4$ (≈ -6 dB). The fact that both minima are very low makes it clear that permutation coding generally requires additional mechanisms for embedding distortion control, which will be dealt with in Section V-B.

Finally, in view of our discussion about the analogy between (5) and (6), it is trivial to see why the average embedding distortion of capacity-achieving perfect steganography for zero-mean X in [6] coincides with $\underline{\xi}$ for zero-sum $\mathbf{x}$, that is to say, $\underline{\xi} \approx -3$ dB.

b) Asymptotics: We will study next the asymptotic behavior for large n of the power a histogram-preserving watermark drawn uniformly at random, corresponding to the encoder choosing messages uniformly at random. Our aim is to quantify a condition under which (3) is a good predictor of the power of *any* histogram-preserving watermark, via the weak law of large numbers. In order to do so we will obtain Chebyshev's inequality for the random variable $\|\boldsymbol{W}\|^2 = 2(\|\mathbf{x}\|^2 - \mathbf{x}^t \boldsymbol{\Pi} \mathbf{x})$, defined using (2) and assuming that $\boldsymbol{\Pi}$ is a random variable uniformly distributed over the ensemble of all permutation matrices. We know already that $\mathrm{E}\{\|\boldsymbol{W}\|^2\} = \overline{\|\mathbf{w}\|^2}$, and therefore we just need to obtain the second moment of $\|\boldsymbol{W}\|^2$ in order to compute its variance. As $\mathbf{x}^t \Pi \mathbf{x} = \mathbf{x}^t \Pi^t \mathbf{x}$, this moment can be put as

$$\mathrm{E}\big\{\|\boldsymbol{W}\|^4\big\} = 4\left(\|\mathbf{x}\|^4 - 2\|\mathbf{x}\|^2 \mathbf{x}^t \,\mathrm{E}\{\boldsymbol{\Pi}\}\mathbf{x} + \mathbf{x}^t \,\mathrm{E}\big\{\boldsymbol{\Pi}\mathbf{x}\mathbf{x}^t\boldsymbol{\Pi}^t\big\}\,\mathbf{x}\right),$$

and hence the desired variance is

$$\mathrm{Var}\left\{\|\boldsymbol{W}\|^2\right\} = 4\left(\mathbf{x}^t \,\mathrm{E}\big\{\boldsymbol{\Pi}\mathbf{x}\mathbf{x}^t\boldsymbol{\Pi}^t\big\}\,\mathbf{x} - \left(\mathbf{x}^t \,\mathrm{E}\{\boldsymbol{\Pi}\}\mathbf{x}\right)^2\right). \quad (13)$$

The computation of the two expectations in (13) can be found in the Appendix. Using these results, after some straightforward algebraic manipulations it can be seen that (13) becomes

$$\mathrm{Var}\left\{\|\boldsymbol{W}\|^2\right\} = \frac{1}{n-1}\left(\overline{\|\mathbf{w}\|^2}\right)^2.$$

Finally, we just use $\mathrm{Var}\{\|\boldsymbol{W}\|^2\}$ and $\mathrm{E}\{\|\boldsymbol{W}\|^2\}$ to obtain Chebyshev's inequality for $\|\boldsymbol{W}\|^2$. For any $\gamma > 0$, this yields

$$\Pr\left\{\left|\|\boldsymbol{W}\|^2 - \overline{\|\mathbf{w}\|^2}\right| \geq \gamma\overline{\|\mathbf{w}\|^2}\right\} \leq \frac{1}{\gamma^2(n-1)}. \quad (14)$$

The most relevant feature of this bound is its independence of $\mathbf{x}$. Although Chebyshev's inequality is known to be loose (since it only relies on the first two moments of the distribution) it is also completely general, and it can be read as saying that the embedding distortion associated to a randomly drawn permutation codeword is not likely to be too different from the average $\overline{\|\mathbf{w}\|^2}$ for large n. This fact will be empirically verified in Section VI.

Despite these considerations, one might still be concerned about the rare instances in which $\|\boldsymbol{W}\|^2$ clearly deviates from the average. We will see next that the geometry of permutation coding strictly confines the maximum distortion (8) in terms of the average distortion (3).

c) Geometry: As noted by Slepian [12], the two basic observations to be made about the geometry of permutation codes are: 1) since $\|\mathbf{y}\| = \|\mathbf{x}\|$, then all codewords lie on an n-dimensional permutation sphere centered at $\mathbf{0}$ with radius $\|\mathbf{x}\|$; and 2) the codewords are really $n-1$ dimensional, as they also lie on the permutation plane with equation $\mathbf{y}^t\mathbf{1} = \mathbf{x}^t\mathbf{1}$.

As we will show next, relevant geometric insights for the embedding distortion analysis can be obtained from the covering sphere of the permutation code. This is a sphere with center $\mathbf{c} \in \mathbb{R}^n$ and minimum radius R_c such that for any codeword $\mathbf{y}$ it holds that $\|\mathbf{y} - \mathbf{c}\|^2 \leq R_c^2$. Since the intersection of the permutation sphere and the permutation plane is a sphere in $n-1$ dimensions that contains all histogram-preserving codewords, then this intersection must also be the intersection of the covering sphere and the permutation plane, and all codewords must be equidistant from the center of the covering sphere. In order to obtain $\mathbf{c}$ and R_c we start by computing the average of all codewords $\overline{\mathbf{y}} \triangleq (1/r)\sum_{m=1}^{r} \mathbf{y}^{(m)} = (1/r)\sum_{\boldsymbol{\sigma}\in\mathcal{S}_{\mathbf{x}}} \Pi_{\boldsymbol{\sigma}}\mathbf{x}$. Using (68) and (69) it can be seen that this average vector is

$$\overline{\mathbf{y}} = \frac{1}{n}(\mathbf{x}^t\mathbf{1})\mathbf{1}. \quad (15)$$

So $\overline{\mathbf{y}}$ is a point with all its coordinates equal to the average of $\mathbf{x}$. As indicated by $\overline{\mathbf{y}}^t\mathbf{1} = \mathbf{x}^t\mathbf{1}$, $\overline{\mathbf{y}}$ lies on the permutation plane. Now, the square of the Euclidean distance of an arbitrary codeword $\mathbf{y}$ to $\overline{\mathbf{y}}$ is

$$\|\mathbf{y} - \overline{\mathbf{y}}\|^2 = \|\mathbf{x}\|^2 - \|\overline{\mathbf{y}}\|^2, \quad (16)$$

where we have used $\mathbf{y}^t\mathbf{1} = \mathbf{x}^t\mathbf{1}$ and $\mathbf{y}^t\overline{\mathbf{y}} = (1/n)(\mathbf{x}^t\mathbf{1})^2 = \|\overline{\mathbf{y}}\|^2$. As (16) is independent of $\mathbf{y}$, then it must also be the square of the covering radius, R_c^2, and the center of the covering sphere must then be $\mathbf{c} = \overline{\mathbf{y}}$. Observing (3) we can thus write

$$R_c^2 = \frac{1}{2}\overline{\|\mathbf{w}\|^2}. \quad (17)$$

Therefore when $\mathbf{x}^t\mathbf{1} \neq 0$ all codewords lie simultaneously on two different spheres: the covering sphere centered at $\overline{\mathbf{y}}$ with radius R_c, and the permutation sphere. The equation of the plane where these two spheres intersect must be the permutation plane $\mathbf{y}^t\mathbf{1} = n\sqrt{\|\mathbf{x}\|^2 - R_c^2} = \mathbf{x}^t\mathbf{1}$. Obviously the covering radius cannot be greater than the radius of the permutation sphere, as it can be seen, for example, from (17) and (4). Therefore

$$R_c \leq \|\mathbf{x}\|, \quad (18)$$

with equality when $\mathbf{x}^t\mathbf{1} = 0$. This is the reason why the covering sphere was never considered in previous works devoted to the application of permutation codes to channel coding [12] or source coding [18], since in those scenarios $\mathbf{x}^t\mathbf{1} = 0$ is usually necessary for energy minimization (and also feasible, since $\mathbf{x}$ is chosen in both), and hence the two spheres coincide.

Using the triangle inequality we can verify next that

$$\|\mathbf{w}\| = \|(\mathbf{y} - \overline{\mathbf{y}}) - (\mathbf{x} - \overline{\mathbf{y}})\| \leq 2\|\mathbf{y} - \overline{\mathbf{y}}\| = 2R_c, \quad (19)$$

or, equivalently, that $\|\mathbf{w}\|$ cannot be greater than the diameter of the covering sphere. Using (17), we see that (19) implies the inequality

$$\|\mathbf{w}\|_{\max}^2 \leq 2\overline{\|\mathbf{w}\|^2}, \quad (20)$$

which complements (10). Trivially, a case in which equality holds in (20) is $\mathbf{x} = v\mathbf{1}$, but there may also exist other equality solutions of (20) with a nonconstant host. Inequality (20) can also be used to narrow down the bound on the probability in (14), since it it implies that this probability can only be nonzero when $\gamma \in (0, 1]$. Finally, when (20) is normalized by the power of the host we obtain

$$\xi_{\min} \geq \underline{\xi}/2, \quad (21)$$

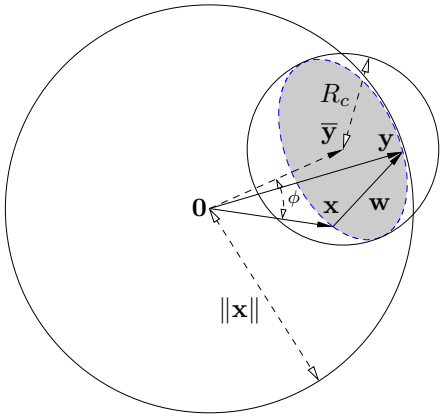

Fig. 1. Schematic of the geometry of first-order perfect steganography.

which means that the host to maximum watermark power ratio is, in decibels, always greater or equal than the host to average watermark power ratio minus approximately 3 dB. This important property will still hold when embedding distortion control is implemented in Section V-B.

An alternative proof of inequality (9) is found, for instance, by combining (19) and (18). This can also be seen directly by applying the triangle inequality as in (19) but with $\mathbf{0}$ instead of $\overline{\mathbf{y}}$, which is equivalent to saying that $\|\mathbf{w}\|$ cannot be greater than the diameter of the permutation sphere. Thus (9) is met with equality whenever there exist two antipodal codewords, that is to say, when there exist $\mathbf{y}$ and $\mathbf{y}'$ such that $\mathbf{y}' = -\mathbf{y}$. This happens when $\overrightarrow{\mathbf{x}}$ is such that $\overrightarrow{x}_i = -\overrightarrow{x}_{n-i+1}$ for $i = 1, 2, \ldots, n$, the antipodal codewords being $\mathbf{y} = \overrightarrow{\mathbf{x}}$ and $\mathbf{y}' = \overleftarrow{\mathbf{x}}$. Also, in the special case in which $\mathbf{x}$ lies in the nonnegative (or nonpositive) orthant, the greatest possible diameter of the covering sphere allows us to replace (9) by $\|\mathbf{w}\|^2_{\max} \leq 2\|\mathbf{x}\|^2$ (and therefore $\xi_{\min} \geq 1/2$), with equality when at least half of the elements of $\mathbf{x}$ are zero.

To conclude this section, it is also interesting to observe that the host to average watermark power ratio can be expressed as a single function of the angle ϕ between $\mathbf{x}$ and $\mathbf{1}$ (equivalently, between any codeword $\mathbf{y}$ and $\mathbf{1}$). Since $\cos\phi = \mathbf{x}^t\mathbf{1}/\left(\|\mathbf{x}\|\|\mathbf{1}\|\right)$, it can be seen from (3) and (11) that

$$\underline{\xi} = \frac{1}{2\sin^2\phi}. \tag{22}$$

This expression also allows us to establish (21) without explicitly resorting to the covering sphere. Since the angle 2ϕ is the opening angle of the right cone with apex $\mathbf{0}$ and base the intersection of the permutation sphere and the permutation plane, then the maximum distance between any codeword $\mathbf{y}$ and $\mathbf{x}$ is bounded as follows:

$$\|\mathbf{y} - \mathbf{x}\|^2 \leq 2\|\mathbf{x}\|^2(1 - \cos 2\phi). \tag{23}$$

Using next the trigonometric identity $\cos 2\phi = 1 - 2\sin^2\phi$ and (22) in (23) we recover (21).

Some of the facts discussed in this section are schematically illustrated in Figure 1.

2) Degree of host change: The degree of host change is an embedding distortion measure alternative to the squared Euclidean distance, which, as we will see, is also insightful in a number of ways. The degree of host change associated to the use of a codeword $\mathbf{y}$ is defined as $\nu \triangleq (1/n)\,\delta(\mathbf{y}, \mathbf{x})$, that is to say, as the Hamming distance per symbol between $\mathbf{y}$ and $\mathbf{x}$, or as the fraction of elements of $\mathbf{y}$ which differ from the same index elements in $\mathbf{x}$.

In order to determine the average degree of host change ($\overline{\nu}$) we firstly define an auxiliary $q \times n$ matrix Λ whose entries are $(\Lambda)_{k,i} = \mathbb{1}_{\{v_k = x_i\}}$, and we let $\Omega \triangleq \Lambda^t\Lambda$. Now, $\operatorname{tr}\Lambda\Pi_{\boldsymbol{\sigma}}\Lambda^t = \operatorname{tr}\Omega\Pi_{\boldsymbol{\sigma}}$ is the number of elements in $\mathbf{y} = \Pi_{\boldsymbol{\sigma}}\mathbf{x}$ unchanged with respect to $\mathbf{x}$. Therefore $\delta(\mathbf{x}, \mathbf{y}) = n - \operatorname{tr}\Omega\Pi_{\boldsymbol{\sigma}}$, and, when the messages are uniform, the average degree of host change can be put as

$$\overline{\nu} = \frac{1}{r}\sum_{\boldsymbol{\sigma}\in\mathcal{S}_{\mathbf{x}}}\frac{1}{n}(n - \operatorname{tr}\Omega\Pi_{\boldsymbol{\sigma}}). \tag{24}$$

We can develop this expression using the equality $(1/r)\sum_{\boldsymbol{\sigma}\in\mathcal{S}_{\mathbf{x}}}\operatorname{tr}\Omega\Pi_{\boldsymbol{\sigma}} = (1/n!)\sum_{\boldsymbol{\sigma}\in\mathcal{S}_n}\operatorname{tr}\Omega\Pi_{\boldsymbol{\sigma}}$, which holds because the second summation contains the same summands as the first one, but each of them repeated $h_1!h_2!\cdots h_q!$ times. As the trace operator is linear, using equation (68) and $\operatorname{tr}\Omega\mathbf{1}\mathbf{1}^t = \mathbf{1}^t\Omega\mathbf{1} = \|\Lambda\mathbf{1}\|^2 = \|\mathbf{h}\|^2$, it is straightforward to see that

$$\overline{\nu} = 1 - \frac{\|\mathbf{h}\|^2}{n^2} = 1 - \|\mathbf{p}\|^2, \tag{25}$$

where $\mathbf{p} \triangleq (1/n)\mathbf{h}$ is the type of $\mathbf{x}$. As $\mathbf{p}$ contains the probabilities of a pmf (whose support is $\mathbf{v}$) then $1/q \leq \|\mathbf{p}\|^2 \leq 1$, and so we have that

$$0 \leq \overline{\nu} \leq 1 - \frac{1}{q}. \tag{26}$$

A useful probabilistic interpretation of (25) can be obtained as follows. Consider two independent discrete random variables whose distribution is the type of $\mathbf{x}$, which we denote as $X \sim \mathbf{p}$ and $Y \sim \mathbf{p}$. The complement of their *index of coincidence*, or, equivalently, the probability of drawing a different outcome in two independent trials of X and Y, is $\Pr\{X \neq Y\} = \sum_{v\in\mathcal{V}} p(X = v)p(Y \neq v)$. This amount is

$$\Pr\{X \neq Y\} = \sum_{k=1}^{q} p_k\,(1 - p_k) = \overline{\nu}. \tag{27}$$

We should also discuss the relationship between $\overline{\nu}$ and $\overline{\|\mathbf{w}\|^2}$. Since $\delta(\mathbf{x}, \mathbf{y}) \leq \|\mathbf{x} - \mathbf{y}\|^2$ for $\mathbf{x}, \mathbf{y} \in \mathbb{Z}^n$ then

$$\overline{\nu} \leq \frac{1}{n}\overline{\|\mathbf{w}\|^2}. \tag{28}$$

Intuitively, a high $\underline{\xi}$ implies a small $\overline{\nu}$. This trivial bound can be greater than one, but it is tight when $\underline{\xi} \to \infty$. Finally note that, unlike the ξ ratios, $\overline{\nu}$ is not preserved by unitary linear transforms.

B. Embedding Rate

In this section we analyze the embedding rate of permutation coding, and we establish bounds connecting this rate to the distortion measures studied in Section IV-A. We also discuss the so-called embedding efficiency, a performance measure which involves both embedding rate and embedding distortion.

The embedding rate associated to the permutation code based on host $\mathbf{x}$ is $\rho \triangleq (1/n)\log r$ bits/host element. Trivially, this rate is necessarily optimum for first-order perfect steganography. In order to obtain a probabilistic perspective of this quantity, assume next that Stirling's approximation $\log_e z! \approx z\log_e z - z$ (for large z) holds for all factorials in (1). In this case the rate can be informally approximated as

$$\rho \approx -\sum_{k=1}^{q} \frac{h_k}{n} \log \frac{h_k}{n} \quad \text{bits/host element.}$$

If we define X to be a discrete random variable whose probability mass function is the type of $\mathbf{x}$, $X \sim \mathbf{p}$, then

$$\rho \approx H(X). \tag{29}$$

In the context of permutation coding, this interpretation of the multinomial coefficient in terms of the entropy of the subjacent type was first given by Berger et al. [18]. More rigorously, approximation (29) is supported by the following bounds (see for instance [20]):

$$H(X) - \zeta^{(q,n)} \leq \rho \leq H(X), \tag{30}$$

where $\zeta^{(q,n)} \triangleq (q/n)\log(n+1)$, which show that $\rho \to H(X)$ as $n \to \infty$, because in this case we have that $\zeta^{(q,n)} \to 0$. We denote this asymptotic equality as $\rho \doteq H(X)$. The upper bound in (30) on the capacity of perfect steganography was previously given by Cachin [2] using the method of types, and by Comesaña and Pérez-González [6] and Wang and Moulin [7] departing from Gel'fand and Pinskers' capacity formula, whereas Ryabko and Ryabko [11] gave a lower bound alternative to the one in (30). Craver and Yu [21] also reasoned along the lines in this section to argue that the so-called "square root law" of steganographic capacity [22] must be wrong.

1) Rate-distortion bounds: Ideally one would like to have exact rate-distortion relationships, that is to say, explicit or implicit equations relating the embedding rate ρ and either $(1/n)\overline{\|\mathbf{w}\|^2}$ or $\overline{\nu}$. Leaving aside the asymptotic case of the binary Hamming setting discussed in Section IV-C1, we suspect that, in general, such relationships do not exist. If they would, they could not simply involve the three aforementioned amounts. Two particular nonbinary hosts validating this observation are $\mathbf{x} = [1,2,3,4,4,4,4]^t$ and $\mathbf{x}' = [1,1,2,2,3,3,3]^t$, which have the same ρ but different $(1/n)\overline{\|\mathbf{w}\|^2}$ and $\overline{\nu}$.

Nevertheless, a general upper bound on ρ solely based on $(1/n)\overline{\|\mathbf{w}\|^2}$ is possible using the differential entropy upper bound on discrete entropy, independently found by Djackov [23], Massey [24], and Willems (unpublished, see [20, Problem 8.7]). This bound is $H(X) < (1/2)\log(2\pi e(\sigma_X^2 + 1/12))$ for a discrete random variable X with support set $\mathcal{V} \subseteq \mathbb{Z}$ and variance σ_X^2. Using the upper bound in (30), since $X \sim \mathbf{p}$ has support set $\mathcal{V} \subset \mathbb{Z}$ and variance $\sigma_{\mathbf{x}}^2$ [given in (5)], we have that the Djackov-Massey-Willems bound yields

$$\rho < \rho_u \triangleq \frac{1}{2}\log\left(2\pi e\left(\frac{\overline{\|\mathbf{w}\|^2}}{2n} + \frac{1}{12}\right)\right). \tag{31}$$

The bound is loose as $\overline{\|\mathbf{w}\|^2}/n \to 0$, as in this case $\rho_u \to 0.2546$ whereas we know that $\rho \to 0$. In practical terms, this effect only starts to be discernible when $\underline{\xi}^* \gtrapprox 50$ dB for $b = 8$.

Furthermore, it is possible to obtain two lower bounds on ρ based on $\overline{\nu}$. The first one uses the inequality $\Pr\{X = Y\} \geq 2^{-H(X)}$ with X and Y i.i.d. [20, Lemma 2.10.1], which is due to the fact that the Rényi entropy bounds the Shannon entropy from below. Using the complement of the index of coincidence (27) and the lower bound in (30) in this inequality, we obtain

$$\rho \geq \rho_l \triangleq -\log(1-\overline{\nu}) - \zeta^{(q,n)}. \tag{32}$$

Asymptotically, we have that $\rho_l \doteq -\log(1-\overline{\nu})$. Since the original bound is sharp when X is uniform, asymptotic equality is achieved in inequality (32) when $\mathbf{h} = (n/q)\mathbf{1}$, and in this case $\rho = \rho_l \doteq \log q$, which corresponds to the maximum average degree of host change $\overline{\nu} = 1 - 1/q$.

The second lower bound is obtained by using the inequality $H(X) \geq 2\Pr\{X \neq Y\}$ with X and Y i.i.d., found by Harremoës and Topsøe [25, Theorem II.6 with $k = 1$]. Using again the lower bound in (30) and (27) we obtain

$$\rho \geq \rho_l' \triangleq 2\,\overline{\nu} - \zeta^{(q,n)}, \tag{33}$$

which is sharper than (32) when $0 < \overline{\nu} < 1/2$. The asymptotic bound is now $\rho_l' \doteq 2\overline{\nu}$.

Finally, from (31-33) we have two upper bounds on $\overline{\nu}$ based on $(1/n)\overline{\|\mathbf{w}\|^2}$ in addition to (28). If we define $\tau \triangleq 2^{\zeta^{(q,n)}}(2\pi e(\overline{\|\mathbf{w}\|^2}/2n + 1/12))^{1/2}$, then from $\rho_l < \rho_u$ and from $\rho_l' < \rho_u$ we respectively have that

$$\overline{\nu} < 1 - \frac{1}{\tau}, \quad \text{and} \quad \overline{\nu} < \frac{1}{2}\log\tau. \tag{34}$$

Interestingly, the first upper bound in (34) cannot be greater than the unity, unlike the second upper bound or (28). However, as $\underline{\xi} \to \infty$ (28) is eventually tighter than both inequalities in (34), mainly due to the lack of asymptotic sharpness of (31).

Since ρ, $\overline{\|\mathbf{w}\|^2}$ and $\overline{\nu}$ are completely determined by $\mathbf{x}$, the rate-distortion bounds presented in this section may look at this juncture like little more than a curiosity. However their relevance will become apparent when we address embedding distortion control for permutation coding in Section V-B.

2) Embedding efficiency: We will study next the average embedding efficiency ($\overline{\varepsilon}$) [26]. This quantity is defined as the average number of message bits embedded per host element change, and, hence, it simultaneously involves embedding rate and embedding distortion. The original idea behind the definition of $\overline{\varepsilon}$ was having a measure of the security of a steganographic algorithm. The rationale was as follows: given two algorithms with the same ρ, the one with higher $\overline{\varepsilon}$ should arguably be more secure since, on average, it would embed the same amount of information with less degree of host change. Loosely speaking, a smaller degree of host change can be related to a smaller probability of detection by the warden, although in fact the last word on this matter can only be had by an optimum detection test.

Since permutation coding leads to perfect steganography with finite memoryless hosts, it may seem that there is little point in contemplating $\overline{\varepsilon}$ as a security measure here. However, we will see that the average embedding efficiency allows for an insightful comparison between permutation coding and model-based steganography [27]. Moreover, $\overline{\varepsilon}$ may also find

application in realistic scenarios in which the memoryless assumption is only an approximation. The computation of $\overline{\varepsilon}$ will require again the Ω matrix defined in Section IV-A2. Firstly see that the embedding efficiency for the message encoded by $\mathbf{y} = \Pi_{\boldsymbol{\sigma}}\mathbf{x}$ is $\varepsilon \triangleq \log r/(n\nu) = \log r/(n - \operatorname{tr}\Omega\Pi_{\boldsymbol{\sigma}})$ bits/host element change. This amount is infinite for $\boldsymbol{\sigma}_0 \in \mathcal{S}_{\mathbf{x}}$ such that $\Pi_{\boldsymbol{\sigma}_0}\mathbf{x} = \mathbf{x}$, which without loss of generality may be assumed to be $\boldsymbol{\sigma}_0 \triangleq [1, 2, \ldots, n]^t$. In order to sidestep this infinity we will consider that the embedding efficiency is zero when $\mathbf{y} = \mathbf{x}$. Therefore when all messages are equally likely the average sought is obtained as

$$\overline{\varepsilon} = \frac{1}{r} \sum_{\boldsymbol{\sigma} \in \mathcal{S}_{\mathbf{x}} \backslash \boldsymbol{\sigma}_0} \frac{\log r}{n - \operatorname{tr}\Omega\Pi_{\boldsymbol{\sigma}}} \quad \text{bits/host element change.} \tag{35}$$

An exact evaluation of (35) requires enumerating the number of permutations associated to each possible value of $\operatorname{tr}\Omega\Pi_{\boldsymbol{\sigma}}$ for $\boldsymbol{\sigma} \in \mathcal{S}_{\mathbf{x}} \backslash \boldsymbol{\sigma}_0$, namely $\{0, 1, 2, \ldots, n-2\}$ ($n-1$ is not a possible value because an elementary permutation involves swapping two indices). It is interesting to note that obtaining this enumeration is equivalent to solving a generalization of the classic problem of *rencontres* (coincidences, or partial derangements) [28], which asks for the number permutations in $\mathcal{S}_n$ that exhibit a given number of fixed points with respect to $\boldsymbol{\sigma}_0$. The generalized problem at hand requires instead finding the number of rearrangements of $\mathbf{x}$ that exhibit a given number of fixed points with respect to $\mathbf{x}$. Whereas the solution to the classic problem is known for arbitrary n (see for instance [29, example 4.2.1]) we are not aware of a solution to the generalized problem for arbitrary $\mathbf{x}$, although an expression for the number of derangements (rearrangements with no fixed points) is known [30]. However, a simple and useful lower bound on $\overline{\varepsilon}$ can be found by observing that (35) involves the harmonic mean of $r-1$ positive values, which is bounded from above by their arithmetic mean [19]. Then we have that

$$\overline{\varepsilon} \geq \overline{\varepsilon}_l \triangleq n\rho \left(\frac{r-1}{r}\right) \left(\frac{1}{r-1} \sum_{\boldsymbol{\sigma} \in \mathcal{S}_{\mathbf{x}} \backslash \boldsymbol{\sigma}_0} (n - \operatorname{tr}\Omega\Pi_{\boldsymbol{\sigma}})\right)^{-1}. \tag{36}$$

As the sum over $\boldsymbol{\sigma} \in \mathcal{S}_{\mathbf{x}} \backslash \boldsymbol{\sigma}_0$ in (36) is equal to the same sum over $\boldsymbol{\sigma} \in \mathcal{S}_{\mathbf{x}}$ it is straightforward to see using (24) that

$$\overline{\varepsilon}_l = \left(\frac{r-1}{r}\right)^2 \frac{\rho}{\overline{\nu}}. \tag{37}$$

Recalling the lower bound in (30) we can in turn bound (37) from below as follows: $\overline{\varepsilon}_l \geq \overline{\varepsilon}'_l \triangleq ((r-1)/r)^2 (H(X) - \zeta^{(q,n)})/\overline{\nu}$. In the following we will consider $\overline{\varepsilon}'_l$ asymptotically as $n \to \infty$. As in this case $(r-1)/r \to 1$ we have that

$$\overline{\varepsilon}'_l \doteq \frac{-\sum_{k=1}^{q} p_k \log p_k}{1 - \sum_{k=1}^{q} p_k^2}. \tag{38}$$

A trivial but loose lower bound on (38) can be found by applying the well-known inequality $p - 1 \geq \log_e p$ to every logarithm in the numerator of the expression, which yields $\overline{\varepsilon}'_l \dot{>} 1/(\log_e 2) \approx 1.44$, but the sharpest lower bound is obtained by applying to (38) the same inequality used in (33), which yields the asymptotic lower bound

$$\overline{\varepsilon}'_l \dot{\geq} 2. \tag{39}$$

Consequently, a minimum average embedding efficiency of 2 bits/host element change is asymptotically guaranteed when using permutation coding.

C. Binary Host

It is interesting to comment on some of the results in Sections IV-A and IV-B for the special case $q = 2$. This case may arise because the host is intrinsically binary or, as we will see in Section V-B, because we are dealing with a two-valued partition of a nonbinary host, a situation that we will see can appear when implementing embedding distortion control with low embedding distortion constraints. In the binary case the number of rearrangements (1) is simply given by the binomial coefficient $r = \binom{n}{h_2} = n!/\left((n-h_2)!h_2!\right)$, and (38) becomes

$$\overline{\varepsilon}'_l \doteq \frac{-(1-p_2)\log(1-p_2) - p_2 \log p_2}{2(1-p_2)p_2}. \tag{40}$$

Perhaps unexpectedly, this very same expression was previously given by Sallee [5, page 11] for the average embedding efficiency of model-based steganography with step size 2 (by which Sallee means a two-valued partitioning of a host signal using adjacent pairs of histogram bins). Inequality (39) was also mentioned in [5], although the justification therein is, apparently, only empirical and restricted to the binary case —note that our conclusion is based on the aforementioned theoretical result by Harremoës and Topsøe [25], which holds for arbitrary q. Incidentally, the coincidence of Sallee's result with (40) suggests that [5, eq. (4)] is not an exact expectation, but an asymptotic lower bound on that expectation. Obviating this minor technical issue, we will see later that the similarities between permutation coding and model-based steganography, which start with (29) already, do no stop here, and that they are not fortuitous but rather due to underlying principles common to both techniques.

1) Hamming distance: An even more particular binary case is the one which the support set of the host is $\mathcal{V} = \{0, 1\}$. In this case it is quite natural to measure the embedding distortion using the Hamming distance. Interestingly, the squared Euclidean distance between two vectors $\mathbf{x}, \mathbf{y} \in \{0,1\}^n$ happens to be equal to their Hamming distance, that is to say, $\delta(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\|^2$. Consequently the Hamming weight of $\mathbf{x}$, $\omega(\mathbf{x}) \triangleq \delta(\mathbf{x}, \mathbf{0})$, is also equal to its norm, $\omega(\mathbf{x}) = \|\mathbf{x}\|^2$. A further identity that holds true in this case is $\|\mathbf{x}\|^2 = \mathbf{x}^t\mathbf{1}$. As a consequence of these facts, the theoretical analysis can be particularized, and new special results can be found. Firstly see that (3) becomes

$$\overline{\omega(\mathbf{w})} = 2\omega(\mathbf{x}) \left(1 - \frac{1}{n}\omega(\mathbf{x})\right). \tag{41}$$

Undertaking unconstrained maximization of (41) on $\omega(\mathbf{x})$ we have that

$$\overline{\omega(\mathbf{w})} \leq \frac{n}{2} \tag{42}$$

with equality when n is even and $\omega(\mathbf{x}) = n/2$. On the other hand, (8) now takes the form

$$\omega(\mathbf{w})_{\max} = 2\min(\omega(\mathbf{x}), n - \omega(\mathbf{x})). \tag{43}$$

It can also be proved that the Hamming covering sphere provides the same bound on $\omega(\mathbf{w})$ as (43). Using expressions (41) and (43) the figures of merit for the embedding distortion (11) and (12) can be put as a sole function of the Hamming weight per symbol of the host, $\overline{\omega}(\mathbf{x}) \triangleq (1/n)\omega(\mathbf{x})$, as follows:

$$\underline{\xi} = \frac{1}{2(1-\overline{\omega}(\mathbf{x}))}, \tag{44}$$

$$\xi_{\min} = \frac{1}{2}\max\left(1, \frac{\overline{\omega}(\mathbf{x})}{1-\overline{\omega}(\mathbf{x})}\right) \tag{45}$$

whereas $\underline{\xi}^* = 1/\left(2\overline{\omega}(\mathbf{x})(1-\overline{\omega}(\mathbf{x}))\right)$. The inequalities in (12) must still hold; additionally, we have from (42) that $\underline{\xi} \geq 2\overline{\omega}(\mathbf{x})$. Since $\mathbf{x}$ is in the nonnegative orthant, we can see that $\xi_{\min} \geq 1/2$, with equality when $\overline{\omega}(\mathbf{x}) \leq 1/2$ (as discussed in Section IV-A1c). Although (21) must still hold, we can now combine (44) and (45) to obtain the following exact relationship between these two amounts:

$$\xi_{\min} = \frac{1}{2}\max\left(1, 2\underline{\xi} - 1\right).$$

Moreover the average degree of host change is completely equivalent to the distortion measures based on the squared Euclidean distance, since in this case $\overline{\nu} = (1/n)\overline{\omega(\mathbf{w})}$ (see that inserting $p_2 = \overline{\omega}(\mathbf{x})$ in (25) yields (41) without normalisation by n). Thus we can write $\overline{\nu} = 1/\underline{\xi}^* = \overline{\omega}(\mathbf{x})/\underline{\xi}$.

Finally, in the binary Hamming case we have an asymptotically exact relationship between the embedding rate ρ and the average embedding distortion $(1/n)\overline{\omega(\mathbf{w})}$ (or $\overline{\nu}$), since these amounts only depend on $\overline{\omega}(\mathbf{x})$ (only asymptotically in the case of ρ). Thus the rate-distortion bounds in Section IV-B1, although still valid, are unnecessary in this case.

V. Practical Issues

Two fundamental issues arise in the implementation of permutation codes for steganography: a) even for moderate n, the exponentially growing number of rearrangements precludes the implementation of a naive encoding scheme (that is, a lookup table mapping messages to rearrangements of $\mathbf{x}$); and b) in general, a pre-established embedding distortion constraint ensuring semantic fidelity between the information-carrying signal and the host may not be met by a permutation of $\mathbf{x}$ chosen at random (corresponding to encoding a random message), and therefore an embedding distortion control mechanism must be put in place. We discuss these two issues in more detail in the following two sections.

A. Near-Optimum Encoding Algorithm

It can be empirically verified that all nonasymptotic theoretical expressions in Section IV are correct using instances of $\mathbf{x}$ with small n, simply by using lookup tables mapping r messages to the r rearrangements of $\mathbf{x}$. However this naive encoding/decoding approach is unfeasible even for moderately large values of n, let alone for the vector lengths typically found in real-world hosts, due to the exponentially growing size of the lookup table. Therefore an efficient method to encode and decode rearrangements (i.e., messages) is essential for a practical implementation of permutation coding in steganography.

1) Prior art: Laisant [31, pages 180–181] gave the first noniterative way to efficiently and bijectively map elements from $\{1, 2, \ldots, n!\}$ to permutations from $\mathcal{S}_n$, and vice versa (in mathematical terminology, to unrank and rank permutations), by relying on the factorial numeration system. This encoding procedure, which has a $O(n \log n)$ complexity, is now commonly called the Lehmer code [32]. However Laisant's method is not applicable to permutation coding because, in general, each rearrangement of $\mathbf{x}$ corresponds to $h_1!h_2!\cdots h_q!$ different permutations. Dealing with the exact same problem as ourselves, but in the context of the application of permutation codes to source coding with a distortion constraint, Berger et al. [18, Appendix VI] proposed the probably first algorithm to efficiently and bijectively map elements from $\{1, 2, \ldots, r\}$ to rearrangements of $\mathbf{x}$, and vice versa (a similar algorithm was also previously devised by Slepian [33], but without giving explicit details). The complexity of their encoding and decoding algorithm is claimed to be $O(n)$, which would seem to solve the problem. However the algorithm is based on Jelinek's implementation of Shannon-Fano-Elias lossless source coding, and therefore the claim can only be true for small n, as indicated by Pasco [8, page 11]. In spite of this limitation, the algorithm by Berger et al. makes it explicit for the first time the connection between permutation coding and source coding. Finally, in their proposal of near-optimum perfect steganography, Ryabko and Ryabko [11] use an algorithm by the first of these authors for ranking and unranking combinatorial objects with $O((\log n)^c)$ time complexity [34], but the connection of this algorithm with existing optimum source coding methods is not explicit.

2) Arithmetic coding and permutation coding: We will give next an encoding procedure for permutation coding which is implementable with $O(n)$ complexity, and which explicitly uses the duality between optimum perfect universal steganography and optimum lossless universal source coding. First of all, approximation (29) of the embedding rate allows us to make the following informal observation about the aforementioned duality: since all different sequences with the same histogram $\mathbf{h}$ as $\mathbf{x}$ can be represented with $nH(X) \approx \log r$ bits by means of optimum lossless source coding, then optimum lossless source decoding of all different messages of length $\log r$ bits must also be able to deliver all r different sequences with the same histogram $\mathbf{h}$ as $\mathbf{x}$, that is to say, all rearrangements of $\mathbf{x}$.

Arithmetic coding is a near-optimum lossless source coding algorithm with $O(n)$ complexity independently proposed by Pasco [8] and Rissanen [9]. Following our discussion above, this algorithm should suffice to implement permutation coding with the same complexity level. In order to do so, we know that one has to make a dual use of the arithmetic encoder and the arithmetic decoder, but the sole application of this duality does not suffice for finite hosts. For reasons that will become evident soon, one also has to reverse the model adaptation used to compress memoryless sources with arithmetic coding [10], which constitutes an additional aspect of the duality. The standard model adaptation —used to implement source coding— and the reverse model adaptation —proposed here to implement permutation coding— are as follows:

- Source coding: the model is initialized with a count of one for each symbol in the alphabet of the sequence to be encoded. After a symbol is encoded, its count is increased by one. Therefore, after n symbols are encoded the symbol counts contain the histogram values of the symbol sequence already encoded *plus one*.
- Permutation coding: the model is initialized with the histogram values of the n-symbol long sequence to be encoded. After a symbol is encoded, its count is decreased by one —note that zero counts are allowed. Therefore, after all n symbols are encoded all symbol counts contain the zero value.

In order to illustrate the method proposed, assume that we wish to encode $\mathbf{x}$ using adaptive arithmetic coding with the second model adaptation discussed above. The initial right-open interval in arithmetic coding is $\mathcal{I}^{(0)} \leftarrow [0,1)$, and according to the reverse model adaptation we let $\mathbf{h}^{(0)} \leftarrow \mathbf{h}$. Then the i-th stage (for $i = 1, 2, \ldots, n$) comprises the following three steps:

1) Division step: $\mathcal{I}^{(i-1)}$ is exactly divided into nonoverlapping right-open subintervals whose lengths are in proportion to the length of $\mathcal{I}^{(i-1)}$ according to the *nonzero* elements of $\mathbf{h}^{(i-1)}/(\mathbf{1}^t\mathbf{h}^{(i-1)})$. Each subinterval with $h_k^{(i-1)} > 0$ is labeled with symbol v_k from $\mathbf{v}$.
2) Encoding step: the subinterval whose label v_s is equal to x_i is declared to be the next interval $\mathcal{I}^{(i)}$.
3) Adaptation step: let $h_s^{(i-1)} \leftarrow h_s^{(i-1)} - 1$ and then declare $\mathbf{h}^{(i)} \leftarrow \mathbf{h}^{(i-1)}$.

The $nH(X) \approx \log r$ most significant fractional bits shared by the binary representation of the endpoints of the final interval $\mathcal{I}^{(n)}$ constitute the compressed representation of $\mathbf{x}$. This representation can also be put as the most significant $\log r$ fractional bits of the decimal real value $(m'-1)/r \in [0,1)$, for some $m' \in \{1, 2, \ldots, r\}$. Decoding (decompressing) m' back into $\mathbf{x}$ requires the same initialization and the same division and adaptation steps as above, whereas the i-th decoding step simply involves declaring the decoded symbol x_i to be equal to the symbol v_s which labels the subinterval where $(m'-1)/r$ lies. The number m' must be represented with arbitrarily long precision. The main achievement of arithmetic coding with respect to Jelinek's algorithm for Shannon-Fano-Elias coding is implementing this requirement using finite precision arithmetic.

We are now ready to describe the permutation encoder and the permutation decoder. The permutation encoder obtains the rearrangement $\mathbf{y} = e(\mathbf{x}, m)$ by carrying out adaptive arithmetic decoding of m as described above. On the other hand, the permutation decoder retrieves the message m embedded in $\mathbf{y}$, that is, $m = d(\mathbf{y})$, by carrying out adaptive arithmetic encoding of $\mathbf{y}$ as described above. We can now clearly see why the reverse model adaptation is essential: it guarantees that $\mathbf{y} = \Pi_{\boldsymbol{\sigma}}\mathbf{x}$ for some $\boldsymbol{\sigma} \in \mathcal{S}_{\mathbf{x}}$. Crucially, encoder and decoder share $\mathbf{h}$ and $\mathbf{v}$ —the pieces of information required for encoding and decoding— precisely because of this fact. Some words are in order about the optimality of this procedure: $\lceil \log r \rceil$ bits per message would be needed in order to represent all r messages. However if $\lceil \log r \rceil > \log r$ then the encoder will map some different messages to the same rearrangement, which leads to unsolvable ambiguities at the decoder. Therefore, univocal decoding is only guaranteed if $\lfloor \log r \rfloor$ bits are used to represent the messages. This is not a serious limitation because (1) is usually very large in steganographic applications, and then $2^{\lfloor \log r \rfloor} \approx r$. Finally, we would like to remark that this permutation encoding algorithm is closely related to the one proposed by Berger et al. [18], although these authors do not explicitly discuss the connection between the adaptation models in source coding and permutation coding. Furthermore, it was the failure of Jelinek's algorithm —on which Berger et al.'s algorithm relies— to cope with input sequences with arbitrarily large n that spurred the development of arithmetic coding four years after [18] was published.

As we mentioned in the introduction, the connection between steganography and compression was investigated from the very start of steganography research. It was Sallee [5] who first used arithmetic coding to flesh out this connection in model-based steganography, an algorithm that maximizes the steganographic embedding rate under a host model preservation constraint coupled with embedding distortion control. In view of the algorithm presented in this section, the defining difference of permutation coding with respect to model-based steganography is —despite their seemingly unrelated starting points— the use of arithmetic coding with an adaptive empirical model of the host, rather than with a static theoretical model of the host. Note that this close relationship between permutation coding and model-based steganography was anticipated already by the appearance of expressions (29) and (40) in our analysis.

To conclude this section we discuss the implications of the algorithm given in this section for the generation of random variables with an arbitrary distribution from a Bernoulli random variable with probability $1/2$. Howard and Vitter [35] previously realized that, by feeding pseudorandom uniformly distributed messages to an arithmetic decoder it is possible to generate random variables which, asymptotically, exhibit any desired distribution, established by the theoretical probabilistic model used (see also [36]). By feeding pseudorandom uniformly distributed messages to a permutation encoder we are doing just that, with the key difference that the desired distribution is not just achieved asymptotically as $n \to \infty$, but exactly in every n-sequence produced by the encoder.

3) Keyed encoding: A straightforward way for incorporating a symmetric secret key $\boldsymbol{\varsigma}$ into the permutation encoder and the permutation decoder (that is, $e_{\boldsymbol{\varsigma}}(\cdot,\cdot)$ and $d_{\boldsymbol{\varsigma}}(\cdot)$) is by means of the application of a shared secret permutation $\boldsymbol{\varsigma} \in \mathcal{S}_n$ to $\mathbf{y}$ by the encoder, which the decoder then inverts before decoding. A problem with this approach is that the key is not unique. An alternative is using a shared permutation of the vectors $\mathbf{h}$ and $\mathbf{v}$, which are implicitly shared by encoder and decoder. In other words, the algorithm in Section V-A2 stays the same but encoder and decoder use $\mathbf{h}^{\boldsymbol{\varsigma}} = \Pi_{\boldsymbol{\varsigma}}\mathbf{h}$ and $\mathbf{v}^{\boldsymbol{\varsigma}} = \Pi_{\boldsymbol{\varsigma}}\mathbf{v}$ instead of $\mathbf{h}$ and $\mathbf{v}$, where the key is now $\boldsymbol{\varsigma} \in \mathcal{S}_q$. In this case there are $q!$ unique keys because all values in $\mathbf{v}$ are unique, although this is not true in general for $\mathbf{h}$. If we represent the key using $\lfloor \log q! \rfloor$ bits, then $\mathbf{h}^{\boldsymbol{\varsigma}}$ and $\mathbf{v}^{\boldsymbol{\varsigma}}$ can also be found by means of the permutation coding algorithm

that we have given in Section V-A2. Finally, a longer key is possible by using t permutations $\boldsymbol{\varsigma}^{(1)}, \boldsymbol{\varsigma}^{(2)}, \ldots, \boldsymbol{\varsigma}^{(t)} \in \mathcal{S}_q$ and then applying $\boldsymbol{\varsigma}^{((i \bmod t)+1)}$ at the i-th arithmetic coding stage.

B. Partitioned Permutation Coding

As discussed in Section I universal steganography requires enforcing an embedding distortion constraint, but a permutation code based on $\mathbf{x}$ —importantly, over which the encoder has no control— may not meet a pre-established distortion constraint, such as a minimum value of ξ. Critically, ξ can be very low, as discussed in Section IV-A1a. However it can always be raised by restricting the encoder to a judiciously chosen subset of the set of all permutation codewords. In this section we will see how this can be accomplished by means of partitioned permutation coding.

Firstly we will need some definitions. A *partitioning* is defined to be a set of p index vectors $\mathcal{U}_n^p = \{\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_p\}$ such that: 1) $[\mathbf{u}_1^t, \mathbf{u}_2^t, \ldots, \mathbf{u}_p^t]^t \in \mathcal{S}_n$; and 2) $\mathbf{u}_j = \overrightarrow{\mathbf{u}_j}$. Consequently, the lengths $n_1, n_2, \ldots, n_p$ of $\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_p$ fulfill $\sum_{j=1}^p n_j = n$. A partitioning $\mathcal{U}_n^p$ is applied to an n-vector $\mathbf{x}$ by forming p subvectors (partitions) $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_p$ composed by elements of $\mathbf{x}$ indexed by the index vectors. Thus, for $j = 1, 2, \ldots, p$ the elements of subvector $\mathbf{x}_j$ are $(\mathbf{x}_j)_i = (\mathbf{x})_{(\mathbf{u}_j)_i}$ with $i = 1, 2, \ldots, n_j$. The application of partitioning $\mathcal{U}_n^p$ to $\mathbf{x}$ is denoted by $\mathcal{U}_n^p(\mathbf{x}) = \{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_p\}$, which will also be referred to as a partitioning of $\mathbf{x}$.

A partitioning $\mathcal{U}_q^p$ can be thought of as a *support partitioning*. A support partitioning can be applied to the q-vector $\mathbf{v}$ containing the bins of the histogram $\mathbf{h}$ of $\mathbf{x}$, yielding $\mathcal{U}_q^p(\mathbf{v}) = \{\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_p\}$. This partitioning of $\mathbf{v}$ can be used in turn to build a *support-induced partitioning* of $\mathbf{x}$, which we denote by $\mathcal{U}_n^p(\mathcal{U}_q^p, \mathbf{x})$. For $j = 1, 2, \ldots, p$, the index vector $\mathbf{u}_j$ of the support-induced partitioning contains all indices i, with $1 \leq i \leq n$, such that $x_i = (\mathbf{v}_j)_k$, with $1 \leq k \leq n_j$. The subvectors in $\mathcal{U}_q^p(\mathbf{v})$ obviously coincide with the supports of the subvectors in $\mathcal{U}_n^p(\mathcal{U}_q^p, \mathbf{x})$. Therefore, the lengths $q_1, q_2, \ldots, q_p$ of $\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_p$ fulfill $\sum_{j=1}^p q_j = q$ not only for $\mathcal{U}_q^p(\mathbf{v})$ (by definition) but also for $\mathcal{U}_n^p(\mathcal{U}_q^p, \mathbf{x})$. Note that for an arbitrary $\mathcal{U}_n^p(\mathbf{x})$ (i.e., not support-induced) we have that $\sum_{j=1}^p q_j \geq q$, because it is possible that the same value $v \in \mathcal{V}$ may be found in more than one subvector of $\mathbf{x}$.

We can now describe partitioned permutation coding: the encoder uses a partitioning $\mathcal{U}_n^p$ to obtain $\mathcal{U}_n^p(\mathbf{x})$ and then undertakes permutation coding separately on each of the resulting p subvectors; in short, it produces $\mathbf{y}_j = \Pi_{\boldsymbol{\sigma}_j} \mathbf{x}_j$ with $\boldsymbol{\sigma}_j \in \mathcal{S}_{\mathbf{x}_j}$ for $j = 1, 2, \ldots, p$. In practice, $\mathbf{y}_j$ is produced by undertaking adaptive arithmetic decoding of $\lfloor \log r_j \rfloor$ bits of the message to be embedded, relying on the histogram $\mathbf{h}_j$ of $\mathbf{x}_j$ as described in Section V-A2. The elements of vector $\mathbf{y}$ are obtained by piecing together all $\mathbf{y}_j$ subvectors, and so $(\mathbf{y})_{(\mathbf{u}_j)_i} = (\mathbf{y}_j)_i$ with $i = 1, 2, \ldots, n_j$, for $j = 1, 2, \ldots, p$. Trivially, $\mathbf{y}$ still preserves the histogram of $\mathbf{x}$ because it stems from permutations of subvectors of $\mathbf{x}$, and hence $\mathbf{y} = \Pi_{\boldsymbol{\sigma}} \mathbf{x}$ for some $\boldsymbol{\sigma} \in \mathcal{S}_{\mathbf{x}}$. Assuming that the decoder shares the partitioning $\mathcal{U}_n^p$ with the encoder, it firstly obtains $\mathcal{U}_n^p(\mathbf{y})$ and then, for $j = 1, 2, \ldots, p$, it undertakes adaptive arithmetic encoding of subvector $\mathbf{y}_j$ relying on its histogram $\mathbf{h}_j$, thus retrieving all p parts of the message.

Finally observe that if the encoder uses a support partitioning $\mathcal{U}_q^p$ to generate $\mathbf{y}$ from $\mathcal{U}_n^p(\mathcal{U}_q^p, \mathbf{x})$, then the decoder simply has to obtain $\mathcal{U}_n^p(\mathcal{U}_q^p, \mathbf{y})$ in order to decode $\mathbf{y}$. So encoder and decoder may share a support partitioning $\mathcal{U}_q^p$ instead of a partitioning $\mathcal{U}_n^p$.

1) Theoretical analysis of partitioned permutation coding: First of all, we will need to rederive the most relevant results in Sections IV-A and IV-B for partitioned permutation coding using a generic partitioning $\mathcal{U}_n^p$. Partitioned permutation coding generally entails a rate-distortion tradeoff: although it can decrease distortion, it can also decrease the number of embeddable messages. This number is now

$$r = \prod_{j=1}^p r_j, \tag{46}$$

where $r_j = \binom{n_j}{\mathbf{h}_j}$ is the multinomial coefficient associated to subvector $\mathbf{x}_j$. Hence, the theoretical embedding rate $\rho = (1/n) \log r$ can be expressed as

$$\rho = \sum_{j=1}^p \frac{n_j}{n} \rho_j, \tag{47}$$

where $\rho_j = (1/n_j) \log r_j$ is the embedding rate for the j-th partition.

As for the embedding distortion results, firstly see that the average watermark power for partitioned permutation coding is $\overline{\|\mathbf{w}\|^2} = (1/r) \sum_{m_1=1}^{r_1} \cdots \sum_{m_p=1}^{r_p} \sum_{j=1}^p \|\mathbf{w}_j^{(m_j)}\|^2$. This amount can be developed as follows

$$\overline{\|\mathbf{w}\|^2} = \sum_{j=1}^p \frac{1}{r} \left(\prod_{\substack{i=1 \\ i \neq j}}^p r_i \right) \sum_{m_j=1}^{r_j} \|\mathbf{w}_j^{(m_j)}\|^2 = \sum_{j=1}^p \overline{\|\mathbf{w}_j\|^2}, \tag{48}$$

and hence

$$\overline{\|\mathbf{w}\|^2} = 2 \left(\|\mathbf{x}\|^2 - \sum_{j=1}^p \frac{1}{n_j} (\mathbf{x}_j^t \mathbf{1})^2 \right). \tag{49}$$

In parallel to (5), the average watermark power per host element can be put in terms of the sample variances per partition, $\sigma_{\mathbf{x}_j}^2 = \overline{\|\mathbf{w}_j\|^2}/(2n_j)$, as follows:

$$\frac{1}{n} \overline{\|\mathbf{w}\|^2} = 2 \sum_{j=1}^p \frac{n_j}{n} \sigma_{\mathbf{x}_j}^2. \tag{50}$$

It is trivial to see that the maximum watermark power is obtained when $\|\mathbf{w}_j\|^2$ is maximum for all $j = 1, 2, \ldots, p$, and then $\|\mathbf{w}\|_{\max}^2 = \sum_{j=1}^p (\|\mathbf{w}_j\|^2)_{\max}$. Thus, we have that

$$\|\mathbf{w}\|_{\max}^2 = 2 \left(\|\mathbf{x}\|^2 - \sum_{j=1}^p \overrightarrow{\mathbf{x}_j}^t \overleftarrow{\mathbf{x}_j} \right). \tag{51}$$

From expressions (49) and (51) one obtains the figures of merit $\underline{\xi}$ ($\underline{\xi}^*$) and $\xi_{\min}$ for the partitioned problem (see definitions in Section IV-A1). Importantly, inequalities (12) and (21) still hold with partitioned permutation coding, because inequalities (10) and (20) hold for the average and maximum of $\|\mathbf{w}_j\|^2$, for every $j = 1, 2, \ldots, p$.

The number of host changes caused within the j-th partition by a permutation $\boldsymbol{\sigma}_j \in \mathcal{S}_{\mathbf{x}_j}$ is computed using a $q_j \times n_j$ matrix Λ_j defined just like Λ in Section IV-A2 but using subvectors $\mathbf{v}_j$ and $\mathbf{x}_j$. Hence, using the matrices $\Omega_j = \Lambda_j^t \Lambda_j$, for $j = 1, 2, \ldots, p$, we can write the average degree of host change for partitioned permutation coding as $\overline{\nu} = (1/r) \sum_{\boldsymbol{\sigma}_1 \in \mathcal{S}_{\mathbf{x}_1}} \cdots \sum_{\boldsymbol{\sigma}_p \in \mathcal{S}_{\mathbf{x}_p}} (1/n) \left(\sum_{j=1}^{p} n_j - \operatorname{tr} \Omega_j \Pi_{\boldsymbol{\sigma}_j} \right)$. This expression can be developed in the same way as (48) to yield

$$\overline{\nu} = \sum_{j=1}^{p} \frac{1}{r} \left(\prod_{\substack{i=1 \\ i \neq j}}^{p} r_i \right) \sum_{\boldsymbol{\sigma}_j \in \mathcal{S}_{\mathbf{x}_j}} \frac{1}{n} \left(n_j - \operatorname{tr} \Omega_j \Pi_{\boldsymbol{\sigma}_j} \right) = \sum_{j=1}^{p} \frac{n_j}{n} \, \overline{\nu}_j, \tag{52}$$

where $\overline{\nu}_j = 1 - (\|\mathbf{h}_j\|/n_j)^2$.

Finally, the average embedding efficiency for partitioned permutation coding is

$$\overline{\varepsilon} = \frac{1}{r} \sum_{\boldsymbol{\sigma}_1, \ldots, \boldsymbol{\sigma}_p \setminus \boldsymbol{\sigma}_0} \frac{\log r}{\sum_{j=1}^{p} n_j - \operatorname{tr} \Omega_j \Pi_{\boldsymbol{\sigma}_j}}. \tag{53}$$

The first summation in (53) ranges over all $\boldsymbol{\sigma}_j \in \mathcal{S}_{\mathbf{x}_j}$, for $j = 1, 2, \ldots, p$, bar the case $\boldsymbol{\sigma}_0 \in \mathcal{S}_{\mathbf{x}}$ in which $\mathbf{y} = \mathbf{x}$. Using the same bounding strategy as in Section IV-B2, a lower bound to (53) is again (37), where r, ρ and $\overline{\nu}$ are now given by (46), (47) and (52), respectively. Trivially (39) still holds for partitioned permutation coding. So the main results in Section IV-B2 are independent of the partitioning used.

We verify next that the rate-distortion upper bound (31) holds for partitioned permutation coding, independently of the partitioning chosen. Applying (31) individually to the partition rates ρ_j in (47), and then using the concavity of $\log(\cdot)$ and Jensen's inequality, we have that

$$\rho < \sum_{j=1}^{p} \frac{n_j}{n} \, \frac{1}{2} \log \left(2\pi e \, \left(\sigma_{\mathbf{x}_j}^2 + \frac{1}{12} \right) \right) \tag{54}$$

$$\leq \frac{1}{2} \log \left(2\pi e \sum_{j=1}^{p} \frac{n_j}{n} \left(\sigma_{\mathbf{x}_j}^2 + \frac{1}{12} \right) \right). \tag{55}$$

Therefore, using (50) in (55) we recover (31). Note that the second inequality above (Jensen's inequality) is met with equality if and only if $\sigma_{\mathbf{x}_j}^2 = \overline{\|\mathbf{w}\|^2}/(2n)$ for all $j = 1, 2, \ldots, p$.

Furthermore, if we only consider the class of support-induced partitionings, then the rate-distortion lower bounds (32) and (33) also hold for partitioned permutation coding independently of the partitioning chosen. Applying (32) individually to the partition rates ρ_j in (47), and then using the convexity of $-\log(\cdot)$ and Jensen's inequality, we have that

$$\rho \geq \sum_{j=1}^{p} \frac{n_j}{n} \, \left(-\log \left(1 - \overline{\nu}_j \right) - \zeta^{(q_j, n_j)} \right) \tag{56}$$

$$\geq -\log \left(\sum_{j=1}^{p} \frac{n_j}{n} (1 - \overline{\nu}_j) \right) - \frac{q}{n} \log \left(\sum_{j=1}^{p} \frac{q_j}{q} n_j + 1 \right), \tag{57}$$

because $\sum_{j=1}^{p} (q_j/q) = 1$ for support-induced partitionings. Therefore, using (52) and $\sum_{j=1}^{p} (q_j/q) n_j \leq n$ in (57) we recover (32). Following the same steps as in (56) and (57), it is even simpler to show that the rate-distortion lower bound (33) still holds for support-induced partitionings.

To conclude this section we state the *invariance property* of support partitionings, whose relevance in the optimization of partitioned permutation coding will be clear later. Assume that encoder and decoder share a support partitioning $\mathcal{U}_q^p$, and that the decoder possesses *any* rearrangement $\mathbf{y}$ of the host $\mathbf{x}$. Then there will be complete agreement between both parties about all theoretical predictions for partitioned permutation coding: the ones made by the encoder using $\mathcal{U}_n^p(\mathcal{U}_q^p, \mathbf{x})$, and the ones made by the decoder using $\mathcal{U}_n^p(\mathcal{U}_q^p, \mathbf{y})$. This is due to the fact that, since the histogram bin vectors $\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_p$ must be the same in both cases (because $\mathbf{v}$ is identical for $\mathbf{x}$ and $\mathbf{y}$), then $\mathbf{y}_j$ must always be some rearrangement of $\mathbf{x}_j$, for every $j = 1, 2, \ldots, p$. Seeing that all expressions (47–52) are independent of any particular rearrangement of the values within any of the p partitions, and that all rate-distortion bounds (31–33) are independent of any particular support-induced partitioning, then encoder and decoder always find the same theoretical predictions.

We must remark that the invariance property does not hold in general for an arbitrary partitioning $\mathcal{U}_n^p$. In this case, in which the encoder makes theoretical predictions using $\mathcal{U}_n^p(\mathbf{x})$ and the encoder makes theoretical predictions using $\mathcal{U}_n^p(\mathbf{y})$, invariance is only guaranteed when $\mathbf{y}$ has been produced by applying partitioned permutation coding to $\mathbf{x}$ using $\mathcal{U}_n^p$, but not when $\mathbf{y}$ is a generic rearrangement of $\mathbf{x}$.

2) Optimality of partitioned permutation coding: We will present next arguments for the optimality of partitioned permutation coding. We saw in Section IV-B that the embedding rate of unpartitioned permutation coding asymptotically reaches the empirical entropy of the host. In general, partitioned permutation coding will achieve a lower embedding rate. In order to study the achievable rate we need to observe that, as discussed in the introduction, steganography with a distortion constraint is a problem of communications with side information at the encoder.

Let us start with the particularization of Gel'fand and Pinsker's result for the achievable rate of a communications system with noncausal side information at the encoder [37]. Since in steganography with a passive warden the channel is noise-free, then Gel'fand and Pinsker's result can be put as

$$\rho_{\text{gp}}^{*} \triangleq \max_{p(y,u'|x)} \{ I(Y;U') - I(U';X) \} \text{ bits/host element}, \tag{58}$$

where the discrete random variables Y and X, both with support $\mathcal{V}$, represent the information-carrying signal and the host, respectively, and U' is an auxiliary discrete random variable whose support has cardinality $|\mathcal{U}'| \leq 2|\mathcal{V}| + 1 = 2q + 1$. The information-carrying signal Y is a deterministic function of the host signal X and U', i.e., $Y = g(X, U')$, and the maximization in (58) is over all distributions $p(y, u'|x)$. In our problem the maximization must take into account two constraints: 1) an embedding distortion constraint $(1/n) \, \mathrm{E} \left\{ \|\boldsymbol{Y} - \boldsymbol{X}\|^2 \right\} \leq (1/\kappa) \sigma_X^2$ must be enforced; and 2) the distribution of Y must

be identical to the distribution of X (perfect steganography). Barron et al. [38] showed that $Y = g(X, U')$ still holds for the optimum solution under the first constraint (which involves a block of symbols rather than a single symbol). The second constraint allows us to develop the difference of mutual informations in Gel'fand and Pinsker's result as follows:

$$I(Y;U') - I(U';X) = H(X|U') - H(Y|U') \quad (59)$$
$$\leq H(X|U'), \quad (60)$$

where (59) is because $H(Y) = H(X)$ and (60) because discrete entropy is nonnegative. Equality is achieved in (60) when $H(Y|U') = 0$, which happens when $Y = g(U')$, that is, when Y is a deterministic function of U' and independent of X. Then it is sufficient to maximize (60) over all $p(u'|x)$. Note that Wang and Moulin [7, eq. (19)] choose the alternative formulation $I(Y;U') - I(U';X) \leq H(Y|X)$ in their theoretical study of the achievable rate of perfect steganography with a passive warden, and then pose the capacity computation as the maximization of this bound over all $p(y|x)$ such that Y has the same distribution as X within the distortion constraint. This formulation is completely equivalent to (60), but, as we will see next, our bound is more convenient to discuss the optimality of partitioned permutation coding.

Let us next recast the embedding rate of partitioned permutation coding (47) in information-theoretical terms. As in (29), we can now approximate the embedding rate $\rho_j = (1/n_j)\log r_j$ within the j-th partition as $\rho_j \approx H(X|U=j)$, where U is a random variable with support $\{1, 2, \ldots, p\}$ and pmf $p(U=j) = n_j/n$, which reflects the probability that X belongs to the j-th partition, and $X|(U=j)$ is a random variable with support $\mathbf{v}_j$ and pmf $\mathbf{p}_j = (1/n_j)\mathbf{h}_j$. Thus we can approximate (47) as

$$\rho \approx \sum_{j=1}^{p} p(U=j)\; H(X|U=j) = H(X|U). \quad (61)$$

Remarkably, this approximation of the embedding rate of partitioned permutation coding, which is asymptotically exact, has the same mathematical form as the upper bound (60) on Gel'fand and Pinsker's rate maximization functional. The partitioning variable U in (61) is in fact analogous to the variable U' in (60): in the same way that U' determines Y independently of X when equality holds in (60), a partitioning determines y_i nearly independently of x_i (for a given message). We can also draw the parallels between the maximization of (60) and of (61). Conditioned to X, any U' that complies with the distortion constraint determines a feasible rate in (60). In similar manner, given $\mathbf{x}$, any partitioning that complies with the distortion constraint determines a feasible rate in (61). In the coming sections we will address the optimization of partitioned permutation coding subject to the embedding distortion constraint, so that (61) may approach ρ^*_{gp}.

Finally, partitioning plays in this problem the role of binning found in any implementation of coding with noncausal side information. Although nested lattice quantizers are typically used to this end [39], observe that these constructs cannot implement perfect steganography for arbitrary finite hosts. However support partitionings can actually be seen as quantizing the support of the host.

3) Theoretical partitioning optimization: We will start with a discussion of the issues involved in the optimization of partitioned permutation coding under an embedding distortion constraint. In the following we will only consider a maximum constraint on $\overline{\|\mathbf{w}\|^2}$ (equivalently, a minimum constraint on $\underline{\xi}$ or on $\underline{\xi}^*$). The reason for this is two-fold: 1) the average watermark power asymptotically approximates the watermark power associated to any random rearrangement of $\mathbf{x}$ (Section IV-A1b), and it also bounds the maximum power through (21); and 2) from $\overline{\|\mathbf{w}\|^2}$ it is possible to assess the goodness of a partitioning by means of the rate-distortion upper bound (31). Lastly, considering (28) and (34), one sees that a maximum constraint on $\overline{\|\mathbf{w}\|^2}$ implies a maximum constraint on $\overline{\nu}$.

As we know, encoder and decoder can always implement partitioned permutation coding by pre-agreeing a partitioning $\mathcal{U}_n^p$ (or else a support partitioning $\mathcal{U}_q^p$). This strategy will be called *static* partitioning, since it involves no optimization. Nevertheless, we will show in Section V-B4a that a particular static support partitioning can be shown to perform not far away from the rate-distortion bound (31), under a mild assumption on the histogram of the host.

Optimization of partitioned permutation coding requires *adaptive* partitioning: choosing a partitioning $\mathcal{U}_n^p$ to comply with a distortion constraint while maximizing the embedding rate. This is equivalent to solving

$$\rho^* = \max_{\mathcal{U}_n^p(\mathbf{x}):\, \underline{\xi} \geq \kappa} \rho, \quad (62)$$

where ρ and $\underline{\xi}$ are the theoretical embedding distortion and embedding rate corresponding to $\mathcal{U}_n^p(\mathbf{x})$. This is a combinatorial optimization problem, whose solvability will be discussed later. Assume for a moment that the encoder is able to solve (62); it then produces $\mathbf{y}$ using the optimum $(\mathcal{U}_n^p)^*$ found. The decoder needs the very same partitioning before it can proceed to decode the message. The simplest possibility would be that the encoder sends $(\mathcal{U}_n^p)^*$ to the decoder through a side channel. However this style of *nonblind* adaptive partitioning is clearly not an option, since it defeats the whole purpose of steganography.

Therefore we must focus our efforts on *blind* adaptive partitioning, which requires that the decoder be able to obtain the same optimum partitioning as the encoder without knowledge of $\mathbf{x}$, i.e., through the maximization of (62) on $\mathcal{U}_n^p(\mathbf{y})$, rather than on $\mathcal{U}_n^p(\mathbf{x})$. This would be a possible strategy if ρ^* would correspond to a unique $(\mathcal{U}_n^p)^*$. Unfortunately there are no guarantees of uniqueness for arbitrary $\mathbf{x}$. In theory, we can circumvent this problem by limiting ourselves to the class of support-induced partitionings: encoder and decoder may separately evaluate the performance of the partitionings induced by a pre-agreed sequence of all support partitionings, and choose, for instance, the very first one deemed to be optimal. Due to the invariance property of support partitionings, encoder and decoder will agree on the theoretical performance associated to every single support-induced partitioning in the sequence, including the optimum one. Note that arbitrary

partitionings are not amenable to this strategy because they lack the invariance property.

Obviating the fact that the class of support-induced partitionings may not contain the overall optimum in (62), the major issue with this strategy is the size of the space to be sequentially searched. The total number of support partitionings $\mathcal{U}_q^p$ (and thus of support-induced partitionings) is the Bell number $\beta_q \triangleq \sum_{p=1}^{q} \left\{ {q \atop p} \right\}$, as, by definition, the Stirling number of the second kind $\left\{ {q \atop p} \right\}$ gives the number of partitionings of a set of cardinality q into p nonempty partitions. For example, for a host represented using $b = 8$ bits/element, $q = 256$ and we have $\beta_{256} \sim 10^{373}$ possible partitionings[1]. The problem can be somewhat relaxed if we only consider support partitionings which contain connected partitions (that is to say, index vectors with adjacent indices), which makes intuitive sense in terms of distortion minimization. The total number of such partitionings is $\beta'_q \triangleq \sum_{p=1}^{q} \binom{q-1}{p-1} = 2^{q-1}$. This can be seen from the bijection between the partitionings with p nonempty connected partitions and the binary strings of length $q-1$ and Hamming weight $p-1$, which implies that there are $\binom{q-1}{p-1}$ such partitionings. Using again $q = 256$ we now have $\beta'_{256} \sim 10^{76}$ possible partitionings, a noticeably smaller number with respect to the general case but still forbidding.

4) Practical partitioning selection: In spite of the complexity involved in the full optimization of blind adaptive partitioning, we will see next that a suboptimal practical solution to this problem is feasible. The key to an implementable strategy must be a short but representative sequence of support partitionings, in order to avoid traversing a prohibitively long full sequence as discussed in the previous section. Fortunately, a necessary property of partitionings that are candidates to solve the optimization problem (62) can be obtained from the development of the rate-distortion bound (31). Observe that this upper bound is sharpest for partitioned permutation coding when Jensen's inequality in (55) holds with equality, which happens when the sample variance per partition is constant, i.e., when $\sigma^2_{\mathbf{x}_j} = \sigma^2_{\mathbf{x}}$ for all $j = 1, 2, \ldots, p$. Therefore, for a fixed embedding distortion constraint, the embedding rate of a partitioning with uneven sample variance per partition can only approach the upper bound (54) but not (55). Thus constancy of the intrapartition sample variance is a necessary (but obviously not sufficient) condition for a partitioning to be able to achieve ρ_u.

As we will verify in Section VI, for a low embedding distortion constraint the necessary condition can be approximated by means of support partitionings which uniformly divide $\mathbf{v}$ into p connected partitions, which gives a sequence of only q support partitionings. This sequence is denoted as $\{\widetilde{\mathcal{U}}_q^p\}$, where $p = 1, 2, \ldots, q$. The p-th support partitioning in the sequence, $\widetilde{\mathcal{U}}_q^p$, may be obtained by defining the centroids $c_j \triangleq v_1 + (j-1/2)(v_q - v_1)/p$ for $j = 1, 2, \ldots, p$, and then, for $k = 1, 2, \ldots, q$, by assigning index k to index vector $\mathbf{u}_m$ such that $m = \arg\min_{j \in \{1,2,\ldots,p\}} |v_k - c_j|$. Thus the encoder just has to find p between 1 and q such that $\mathcal{U}_n^p(\widetilde{\mathcal{U}}_q^p, \mathbf{x})$ complies with $\underline{\xi} \geq \kappa$ and maximizes ρ, which it uses to produce $\mathbf{y}$. On the other hand the decoder finds p between 1 and q such that $\mathcal{U}_n^p(\widetilde{\mathcal{U}}_q^p, \mathbf{y})$ complies with $\underline{\xi} \geq \kappa$ and maximizes ρ, which it uses to decode $\mathbf{y}$. In case of nonunique solutions, the first one in the sequence, for instance, is chosen by both parties. With respect to the actual performance of this practical approach, it must also be observed that, for a partitioning $\mathcal{U}_n^p(\widetilde{\mathcal{U}}_q^p, \mathbf{x})$, the maximum average degree of host change in each partition is $\overline{\nu}_j = 1 - 1/q_j$, which is achieved when there is uniformity within each of the partitions. This maximizes the overall average degree of host change (52), which in turn maximizes the rate-distortion lower bounds (32) and (33). This condition is in fact approximated by $\mathcal{U}_n^p(\widetilde{\mathcal{U}}_q^p, \mathbf{x})$ as p increases. Since large p generally corresponds to high $\underline{\xi}$, for large κ the rate-distortion lower bounds close in on the upper bound.

For the reasons discussed above (approximate fulfillment of necessary condition for approaching upper bound and approximate maximization of lower bounds), the simple strategy described in this section can achieve embedding rates reasonably close to the upper bound (31), as it will be empirically verified in Section VI. It must be remarked that the upper bound applies to general partitionings (that is, not just to support-induced partitionings), and that it is not attainable with equality.

a) Histogram-preserving LSB steganography: In this concluding section we analyze the performance of partitioned permutation coding using a special static partitioning strategy. Suppose that both encoder and decoder use a static support partitioning with $p = \lceil q/2 \rceil$ connected partitions, each of which contains two adjacent values of the support $\mathbf{v}$, i.e., every subvector $\mathbf{v}_j$ in $\mathcal{U}_q^{\lceil q/2 \rceil}(\mathbf{v})$ contains a pair of adjacent histogram bins $(\mathbf{v}_j)_1$ and $(\mathbf{v}_j)_2 = (\mathbf{v}_j)_1 + 1$. Assuming that the histogram $\mathbf{h}$ varies slowly, it is reasonable to assume that $(\mathbf{h}_j)_1 \approx (\mathbf{h}_j)_2$ for each partition $\mathbf{x}_j$ in $\mathcal{U}_n^{\lceil q/2 \rceil}(\mathcal{U}_q^{\lceil q/2 \rceil}, \mathbf{x})$. The assumption can also be put as $\mathbf{h}_j \approx (n_j/2)\mathbf{1}$, which is equivalent to saying both that there is approximate uniformity within each partition and that the intrapartition variance is approximately constant. To see this last point, it is convenient to write the intrapartition variance in terms of $\mathbf{h}_j$ and $\mathbf{v}_j$ as follows: $\sigma^2_{\mathbf{x}_j} = (1/n_j)(\mathbf{v}_j^t \operatorname{diag}(\mathbf{h}_j)\mathbf{v}_j - (1/n_j)(\mathbf{v}_j^t \mathbf{h}_j)^2)$. Using this expression and the approximation of the histogram within the partition we get

$$\begin{aligned} \sigma^2_{\mathbf{x}_j} &\approx \frac{1}{2}\left(\|\mathbf{v}_j\|^2 - \frac{1}{2}(\mathbf{v}_j^t \mathbf{1})^2 \right) \\ &= \frac{1}{4}\left((\mathbf{v}_j)_2 - (\mathbf{v}_j)_1 \right)^2 = \frac{1}{4}. \end{aligned} \tag{63}$$

We have discussed that the two aforementioned properties are desirable for a good partitioning strategy, which we proceed to verify next. From (49) and (63) we have that the average embedding distortion per host element can be approximated as

$$\frac{1}{n}\overline{\|\mathbf{w}\|^2} \approx \frac{1}{2}. \tag{64}$$

Now, for this embedding distortion the upper bound (31) becomes

$$\rho_u = \frac{1}{2}\log\left(\frac{2}{3}\pi e\right) = 1.2546.$$

[1]When $q = 2$ there is only one support partitioning with nonzero rate, and hence blind adaptive partitioning is not possible for binary hosts.

On the other hand, using the histogram assumption and assuming that n_j is even for simplicity, the embedding rate (47) can be approximated as

$$\rho \approx \sum_{j=1}^{p} \frac{n_j}{n} \left(\frac{1}{n_j} \log \binom{n_j}{\frac{n_j}{2}} \right) \approx 1, \tag{65}$$

where the second approximation assumes that n_j is large for all $j = 1, 2, \ldots, p$ and uses Stirling's formula. Therefore we can see that the embedding rate is approximately 20% below the upper bound with this static partitioning strategy. Even if this looks like a modest achievement, we must bear in mind that (64) and (65) hold for any host for which the slowly varying histogram assumption is valid, whereas ρ_u can only be approached by a host whose histogram resembles a Gaussian distribution.

This analysis also shows that the static partitioning strategy employed in this section can implement histogram-preserving steganography with nearly the same embedding distortion and embedding rate as (non histogram-preserving) LSB steganography —namely (64) and (65). Therefore this partitioning strategy allows partitioned permutation coding to supersede all existing ad-hoc strategies aimed at reducing the histogram artifacts of LSB steganography —not only the histogram-preserving techniques mentioned in Section III, but also the extensively studied LSB "matching" algorithm (also known as ± 1 steganography) [40].

VI. Results

In this section we verify the match between empirical and theoretical results. Each empirical value corresponds to one single message m drawn uniformly at random from $\{0, 1, \ldots, 2^{n\rho_{\text{emp}}}\}$, where the empirical embedding rate is $\rho_{\text{emp}} = (1/n) \sum_{j=1}^{p} \lfloor \log \tilde{r}_j \rfloor$ bits/host element for a partitioning $\mathcal{U}_n^p(\mathbf{x})$. The $\log \tilde{r}_j$ amounts are computed using Robbins' sharpening of Stirling's formula [41]

$$\sqrt{2\pi z}(z/e)^z e^{(12z+1)^{-1}} < z! < \sqrt{2\pi z}(z/e)^z e^{(12z)^{-1}}. \tag{66}$$

The lower bound in (66) is applied to the factorial in the numerator of the multinomial coefficient r_j, and the upper bound in (66) to each of the factorials in its denominator. We thus obtain $\log \tilde{r}_j < \log r_j$. We follow this procedure because it is essential for the near-optimum permutation encoder/decoder that r_j is accurately approximated but never overestimated, because the j-th partition cannot convey more messages than rearrangements of $\mathbf{x}_j$. On the other hand the lower bound in (30), although convenient from a theoretical point of view, is less sharp than the one obtained using Robbins' result. As for the theoretical rate (47), $\log r_j$ is computed by applying the upper bound in (30) to ρ_j. Note that with finite precision arithmetic one can only exactly compute small factorials; when this is possible, no bounds are needed to compute either the empirical or the theoretical embedding rate.

The information-carrying vector $\mathbf{y} = e(\mathbf{x}, m)$ is then produced as described in Section V-B. In every case it is verified that $\mathbf{y}$ preserves the histogram of $\mathbf{x}$, and that the decoder retrieves the message without error, i.e., $d(\mathbf{y}) = m$. The following empirical amounts are computed from $\mathbf{w} = \mathbf{y} - \mathbf{x}$:

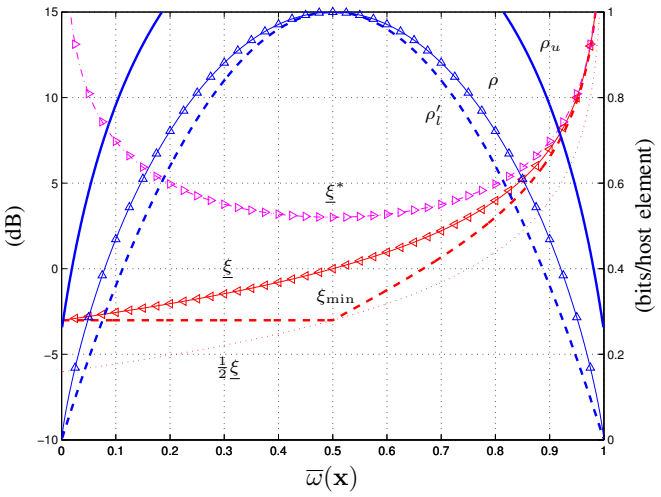

Fig. 2. Performance of unpartitioned permutation coding for a host with support $\mathcal{V} = \{0, 1\}$ and $n = 10^6$ as a function of its Hamming weight per symbol. Lines are theoretical results and symbols are empirical measurements.

$\xi_{\text{emp}} = \|\mathbf{x}\|^2/\|\mathbf{w}\|^2$ (empirical host to watermark power ratio), $\xi^*_{\text{emp}} = n(2^b - 1)^2/\|\mathbf{w}\|^2$ (empirical peak host to watermark power ratio), $\nu_{\text{emp}} = (1/n) \sum_{i=1}^{n} \mathbb{1}_{\{w_i \neq 0\}}$ (empirical degree of host change) and $\varepsilon_{\text{emp}} = \rho_{\text{emp}}/\nu_{\text{emp}}$ bits/host element change (empirical embedding efficiency).

The results in Figure 2 are for unpartitioned permutation coding in the binary Hamming case. As discussed in Section IV-C, in this case the rate-distortion performance only depends on the normalized Hamming weight of the host, and consequently it is particularly easy to visualize. The rate-distortion bounds are unnecessary in this special case, but they are anyway plotted for illustration purposes. In the figure, the lower rate-distortion bound is the asymptotic result $\rho'_l \doteq 2\overline{\nu}$, which is tighter than ρ_l because from (26) we have that $\overline{\nu} \leq 1/2$ for a binary host. On the other hand, the upper bound (31) is not informative for a range of $\overline{\omega}(\mathbf{x})$ around $1/2$, in which it is greater than one. The empirical rate approximates almost exactly the theoretical rate, just because r is very large and the bounds (66) are asymptotically tight. It is more interesting to see that the empirical distortion results —which correspond to single watermarks (i.e., they are not averages)— accurately match the theoretical predictions involving averages, that is, $\underline{\xi}$ and $\underline{\xi}^*$. We showed in Section IV-A1b that this asymptotic behavior of the embedding distortion of permutation coding for large n is a consequence of the law of large numbers. Also see that, as discussed in Section IV-A1c, $\xi_{\min} \geq \underline{\xi}/2$.

We verify next the performance of partitioned permutation coding, using a host drawn from a Gaussian distribution with mean 128 and standard deviation 25, and quantized to $\mathcal{V} = \{0, 1, \ldots, 2^b - 1\}$ ($b = 8$). The sequence of support partitionings $\{\widetilde{\mathcal{U}}_q^p\}$ discussed in Section V-B4 is used to obtain the results in Figure 3. The rate-distortion bounds are now key to the performance achieved by partitioned permutation coding. First of all, like in the previous figure we observe a close match between measurements for single experiments and their corresponding averages. This is also true for the match between the empirical embedding efficiency and the

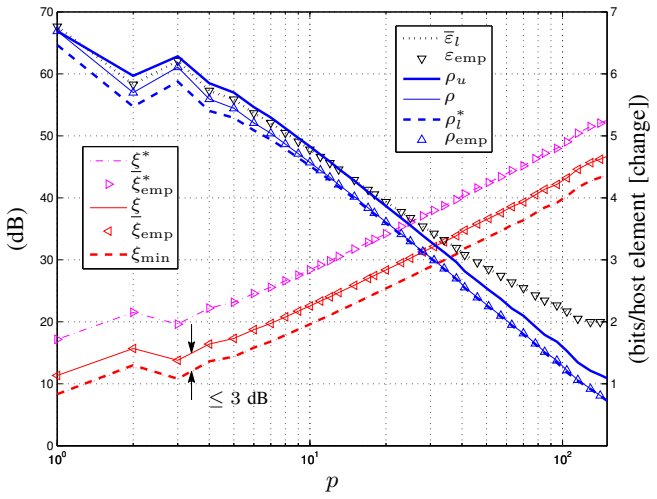

Fig. 3. Performance of partitioned permutation coding for a quantized Gaussian host with $n = 10^6$ as a function of the number of partitions, for the sequence of support partitionings $\{\widetilde{\mathcal{U}}_q^p\}$; $\rho_l^* = \max(\rho_l, \rho_l')$.

lower bound on $\bar{\varepsilon}$, which are also displayed in this case. It can be verified that, as discussed in Section V-B inequalities (21) and (39) still hold for partitioned permutation coding. Finally, the most important feature in Figure 3 is the small constant gap $\rho_u - \rho_{\text{emp}}$ throughout, which is roughly the same as the gap $\rho_u - \rho_l^*$ between the rate-distortion bounds.

VII. Concluding Remarks

We have given a theoretical analysis and a practical solution to the problem of optimum perfect universal steganography of finite memoryless sources with a passive warden. As we have seen, Slepian's Variant I permutation codes are central to this problem. We have shown that these codes can be efficiently implemented by means of adaptive arithmetic coding, which is satisfactory given the long-observed duality between perfect steganography and lossless compression. The procedure to rank and unrank permutations that we have given is general, and may also find application elsewhere. Finally we have studied permutation coding under a distortion constraint, and shown that partitioned permutation coding can perform reasonably close to an unattainable upper bound on the rate-distortion function of the problem.

We expect that the results in this paper may foster the development of well-founded steganographic methods for real-world signals, in the same way that optimum lossless source coding of memoryless signals is at the core of compressing real-world signals. Possible approaches include the use of decorrelating integer-to-integer transforms, or decorrelation through predictive techniques such as prediction by partial matching [42], prior to the application of permutation coding. Inevitably, steganography will no longer be perfect when the memoryless assumption becomes only an approximation, but the aforementioned approaches have the advantage of decoupling the perfect steganography problem from the decorrelation problem. Observe that the breakdown of the memoryless assumption in practical scenarios affects asymmetrically optimum lossless source coding and optimum perfect steganography: in the first problem the perfect recovery of the decompressed signals is not affected, whereas in the second one the undetectability of the steganographic signals is.

Appendix

In this appendix we derive two expectations of matrix functions of a random variable Π uniformly distributed over the ensemble of all permutation matrices, i.e., $\Pr\{\Pi = \Pi_{\boldsymbol{\sigma}}\} = 1/n!$ for all $\boldsymbol{\sigma} \in \mathcal{S}_n$.

A. First Order Expectation

We wish to calculate

$$\mathrm{E}\{\Pi\} = \frac{1}{n!} \sum_{\boldsymbol{\sigma} \in \mathcal{S}_n} \Pi_{\boldsymbol{\sigma}}. \tag{67}$$

In order to undertake this computation consider the number of $n \times n$ permutation matrices which have a one at any given entry. This is equivalent to fixing the corresponding row, and therefore there are $(n-1)!$ possibilities for the remaining $n-1$ rows. Since this holds for any entry, the summation on the right equals $(n-1)!\,\mathbf{1}\mathbf{1}^t$, and so we have that

$$\mathrm{E}\{\Pi\} = \frac{1}{n}\,\mathbf{1}\mathbf{1}^t. \tag{68}$$

Also see that, if we assume $\Pr\{\Pi = \Pi_{\boldsymbol{\sigma}}\} = 1/r$ for all $\boldsymbol{\sigma} \in \mathcal{S}_{\mathbf{x}}$, then the average vector $\mathrm{E}\{\Pi\}\mathbf{x} = \mathrm{E}\{\Pi\mathbf{x}\}$ can be obtained using (68) by observing that

$$\frac{1}{r} \sum_{\boldsymbol{\sigma} \in \mathcal{S}_{\mathbf{x}}} \Pi_{\boldsymbol{\sigma}}\mathbf{x} = \frac{1}{n!} \sum_{\boldsymbol{\sigma} \in \mathcal{S}_n} \Pi_{\boldsymbol{\sigma}}\mathbf{x}, \tag{69}$$

which holds because the second sum of vectors contains the same summands as the first one, but each of them repeated $h_1!h_2!\cdots h_q!$ times.

B. Second Order Expectation

Next we wish to compute

$$\mathrm{E}\left\{\Pi\mathbf{x}\mathbf{x}^t\Pi^t\right\} = \frac{1}{n!} \sum_{\boldsymbol{\sigma} \in \mathcal{S}_n} \Pi_{\boldsymbol{\sigma}}\mathbf{x}\mathbf{x}^t\Pi_{\boldsymbol{\sigma}}^t, \tag{70}$$

for some arbitrary vector $\mathbf{x}$. This average is a particular case of the general formula given by Daniels in [43] to compute $\mathrm{E}\{\Pi\Delta\Pi^t\}$; however, in the case $\Delta = \mathbf{x}\mathbf{x}^t$ of interest to us we can derive the computation in a simpler way by observing that all matrices inside the summation in (70) are of the form $\Pi_{\boldsymbol{\sigma}}\mathbf{x}\,(\Pi_{\boldsymbol{\sigma}}\mathbf{x})^t$. First, see that the diagonal of all these matrices is some permutation of $\mathbf{x}\,\mathrm{diag}(\mathbf{x})$. If we focus on a diagonal entry of one of the matrices and we count the number of matrices for which that entry is x_i^2 we see that there are $(n-1)!$ possibilities, because this amounts to fixing one position in a permutation of the elements of $\mathbf{x}$. Since this holds for any diagonal entry of the matrices within the summation, then all diagonal entries of the matrix sum have the value $(n-1)!\|\mathbf{x}\|^2$. Next, the off-diagonal entries of the matrices in the summation are cross-products of the form x_ix_j, with $i \neq j$. Following the same procedure, if we focus on an off-diagonal entry of one of these matrices and we count the number of matrices for which that entry is x_ix_j, we see that there can only be $(n-2)!$

possibilities, because this amounts to fixing two positions in a permutation of $\mathbf{x}$. Since this holds for any off-diagonal entry of the matrices within the summation, then all off-diagonal entries of the matrix sum are $(n-2)!((\mathbf{x}^t\mathbf{1})^2-\|\mathbf{x}\|^2)$, because $(\mathbf{x}^t\mathbf{1})^2-\|\mathbf{x}\|^2$ is the sum of all x_ix_j cross-products with $i \neq j$. Combining the diagonal and off-diagonal cases we can write

$$\mathrm{E}\left\{\Pi\mathbf{x}\mathbf{x}^t\Pi^t\right\} = a\,\mathrm{I} + b\,\mathbf{1}\mathbf{1}^t,$$

with $b = \left((\mathbf{x}^t\mathbf{1})^2 - \|\mathbf{x}\|^2\right)/(n(n-1))$ and $a+b = (1/n)\|\mathbf{x}\|^2$.

Finally observe that average (70) is the same if obtained over $\mathcal{S}_{\mathbf{x}}$ rather than over $\mathcal{S}_n$, for the same reason that makes expression (69) true.

REFERENCES

[1] P. Moulin and R. Koetter, "Data-hiding codes," *Proc. IEEE*, vol. 93, pp. 2083–2127, December 2005.

[2] C. Cachin, "An information-theoretic model for steganography," in *Procs. of the 2nd Int. Workshop on Information Hiding*, ser. LNCS, vol. 1525. Portland, USA: Springer-Verlag, April 1998, pp. 306–318.

[3] R. Anderson and F. Petitcolas, "On the limits of steganography," *IEEE J. Select. Areas Commun.*, vol. 4, pp. 474–481, May 1998.

[4] C. E. Shannon, "Coding theorems for a discrete source with a fidelity criterion," in *IRE Nat. Conv. Rec., Pt. IV*, 1959, pp. 142–163.

[5] P. Sallee, "Model-based steganography," in *Procs. of the 2nd Int. Workshop on Digital Watermarking (IWDW)*, Seoul, Korea, October 2003, pp. 154–167.

[6] P. Comesaña and F. Pérez-González, "On the capacity of stegosystems," in *Procs. of the 9th ACM Workshop on Multimedia & Security*, Dallas, USA, September 2007, pp. 15–24.

[7] Y. Wang and P. Moulin, "Perfectly secure steganography: Capacity, error exponents, and code constructions," *IEEE Trans. on Inf. Theory*, vol. 54, no. 6, pp. 2706–2722, June 2008.

[8] R. C. Pasco, "Source coding algorithms for fast data compression," Ph.D. dissertation, Stanford University, May 1976.

[9] J. Rissanen, "Generalized Kraft inequality and arithmetic coding," *IBM J. Res. Dev.*, vol. 20, no. 198, 1976.

[10] J. Cleary and I. Witten, "A comparison of enumerative and adaptive codes," *IEEE Trans. Inf. Theory*, vol. 30, no. 2, pp. 306–315, March 1984.

[11] B. Ryabko and D. Ryabko, "Asymptotically optimal perfect steganographic systems," *Problems of Information Transmission*, vol. 45, no. 2, pp. 184–190, February 2009.

[12] D. Slepian, "Permutation modulation," *Procs. of the IEEE*, vol. 53, no. 3, pp. 228–236, 1965.

[13] N. Provos, "Defending against statistical steganalysis," in *Procs. of the 10th USENIX Security Symposium*, Washington D.C., USA, August 2001, pp. 323–335.

[14] E. Franz, "Steganography preserving statistical properties," in *Information Hiding*, ser. Lecture Notes in Computer Science, F. Petitcolas, Ed. Springer, 2003, vol. 2578, pp. 278–294.

[15] M. Kumar and R. Newman, "J3: High payload histogram neutral JPEG steganography," in *8th Annual Int. Conf. on Privacy, Security and Trust (PST)*, Otawa, Canada, August 2010, pp. 46–53.

[16] G. Luo and K. Subbalakshmi, "Zero Kullback-Leibler divergence image data hiding," in *IEEE Global Telecommunications Conf. (GLOBECOM)*, Houston, USA, December 2011, pp. 1–5.

[17] T. Mittelholzer, "An information-theoretic approach to steganography and watermarking," in *Procs. of the 3rd Int. Information Hiding Workshop*, ser. LNCS, vol. 1768. Dresden, Germany: Springer-Verlag, October 1999, pp. 1–16.

[18] T. Berger, F. Jelinek, and J. Wolf, "Permutation codes for sources," *IEEE Trans. Inf. Theory*, vol. 18, no. 1, pp. 160 –169, January 1972.

[19] G. H. Hardy, J. E. Littlewood, and G. Pólya, *Inequalities*. Cambridge at the University Press, 1934.

[20] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. John Wiley & Sons, 2006.

[21] S. Craver and J. Yu, "Subset selection circumvents the square root law," in *Media Forensics and Security II*, ser. SPIE Proceedings, vol. 7541, San José, USA, January 2010.

[22] A. D. Ker, T. Pevný, J. Kodovský, and J. Fridrich, "The square root law of steganographic capacity," in *Proc. 10th ACM Workshop on Multimedia and Security*, 2008, pp. 107–116.

[23] A. G. Djackov, "On a search model of false coins," in *Topics in Information Theory*, ser. Colloquia Mathematica Societatis János Bolyai, I. Csiszár and P. Elias, Eds. Budapest, Hungary: North Holland Publishing Co., 1977, vol. 16, pp. 163–170.

[24] J. Massey, "On the entropy of integer-valued random variables," in *Procs. of Int. Workshop on Information Theory*, Beijing, China, July 1988, pp. C1.2–C1.4.

[25] P. Harremoës and F. Topsøe, "Inequalities between entropy and index of coincidence derived from information diagrams," *IEEE Trans. Inf. Theory*, vol. 47, no. 7, pp. 2944–2960, November 2001.

[26] A. Westfeld, "F5 – A steganographic algorithm," in *Procs. of the 4th Int. Information Hiding Workshop*, ser. LNCS, vol. 2137. Springer-Verlag, 2001, pp. 289–302.

[27] P. Sallee, "Model-based methods for steganography and steganalysis," *Int. Journal of Image and Graphics*, vol. 5, no. 1, pp. 167–190, 2005.

[28] P. Montmort, *Essay d'Analyse sur les Jeux de Hazard*, Paris, 1708.

[29] K. Lange, *Applied Probability*. Springer, 2003.

[30] S. Even and J. Gillis, "Derangements and Laguerre polynomials," *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 79, no. 1, pp. 135–143, January 1976.

[31] C.-A. Laisant, "Sur la numération factorielle, application aux permutations," *Bulletin de la Société Mathématique de France*, vol. 16, pp. 176–183, 1888.

[32] D. Lehmer, "Teaching combinatorial tricks to a computer," *Proc. Sympos. Appl. Math. Combinatorial Analysis, Amer. Math. Soc*, vol. 10, pp. 179–193, 1960.

[33] D. Slepian, "Permutation code signaling," U.S. Patent 3,196,351, July 1965.

[34] B. Ryabko, "Fast enumeration of combinatorial objects," *Diskretnaya Mathematika*, vol. 10, no. 2, pp. 101–119, 1998.

[35] P. Howard and J. Vitter, "Practical implementations of arithmetic coding," in *Image and Text Compression*, J. Storer, Ed. Kluwer Academic Publishers, 1992, pp. 85–112.

[36] D. E. Knuth and A. C. Yao, "The complexity of random number generation," in *Algorithms and Complexity: Recent Results and New Directions*, J. F. Traub, Ed. New York: Academic Press, 1976, pp. 357–428.

[37] S. I. Gel'fand and M. S. Pinsker, "Coding for channel with random parameters," *Problems of Control and Information Theory*, vol. 9, no. 1, pp. 19–31, 1980.

[38] R. J. Barron, B. Chen, and G. W. Wornell, "The duality between information embedding and source coding with side information and some applications," *IEEE Trans. Inf. Theory*, vol. 49, no. 5, pp. 1159–1180, May 2003.

[39] R. Zamir, S. Shamai, and U. Erez, "Nested linear/lattice codes for structured multiterminal binning," *IEEE Trans. Inf. Theory*, vol. 48, no. 6, pp. 1250–1276, June 2002.

[40] T. Sharp, "An implementation of key-based digital signal steganography," in *Procs. of the 4th Information Hiding Workshop*, ser. LNCS, vol. 2137. Springer-Verlag, 2001, pp. 13–26.

[41] H. Robbins, "A remark on Stirling's formula," *The American Mathematical Monthly*, vol. 62, no. 1, pp. 26–29, January 1955.

[42] J. Cleary and I. Witten, "Data compression using adaptive coding and partial string matching," *IEEE Trans. Commun.*, vol. 32, no. 4, pp. 396–402, Apr 1984.

[43] H. E. Daniels, "Processes generating permutation expansions," *Biometrika*, vol. 49, no. 1/2, pp. 139–149, June 1962.